

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
14 March 2002 (14.03.2002)

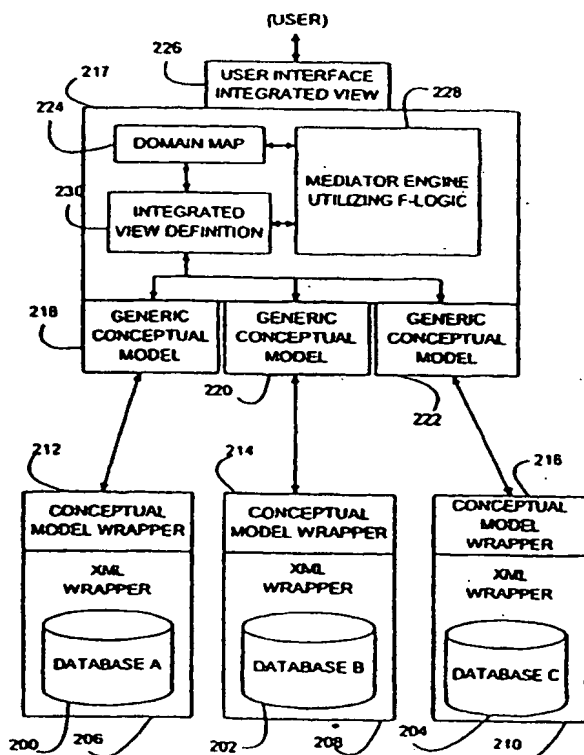
PCT

(10) International Publication Number
WO 02/21259 A1

- (51) International Patent Classification: G06F 7/60, 17/10, 101/00
- (21) International Application Number: PCT/US01/27653
- (22) International Filing Date: 7 September 2001 (07.09.2001)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 60/231,094 8 September 2000 (08.09.2000) US
- (71) Applicant (for all designated States except US): THE REGENTS OF THE UNIVERSITY OF CALIFORNIA [US/US]; 1111 Franklin Street, 5th floor, Oakland, CA 94607-5200 (US).
- (72) Inventors; and
(75) Inventors/Applicants (for US only): GUPTA, Amarnath [IN/US]; 12510 Carmel Creek Road, Apt. 199, San Diego, CA 92130 (US). LUDAESCHER, Bertram [DE/US]; 4128-148 Via Candidiz, San Diego, CA 92130-3160 (US). MARTONE, Maryann, E. [US/US]; 8368G Via Sonoma, La Jolla, CA 92037 (US).
- (74) Agent: FITZSIMONS, Thomas, R.; Greer, Burns & Crain, Ltd., Suite 2500, 300 South Wacker Drive, Chicago, IL 60606 (US).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

[Continued on next page]

(54) Title: DATA SOURCE INTEGRATION SYSTEM AND METHOD



(57) Abstract: A method and program product for integrating different data sources has steps of obtaining semantic information from each of the different data sources (200, 202, 210), creating a conceptual model of (218, 220, 22) the data source using the semantic information, and accessing one or more secondary knowledge sources. The secondary information sources contain information regarding the relations of data from different of the databases, so that an integrated semantic model of all of the databases (200, 202, 210) may be created. Queries can then be processed using the integrated semantic model.

WO 02/21259 A1

BEST AVAILABLE COPY



(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

— with international search report

DATA SOURCE INTEGRATION SYSTEM AND METHOD

Technical Field

The present invention is related to integrating information from a plurality of different information sources. More particularly, the present invention is related to methods and computer program products for creating an integrated model that correlates information from a plurality of different heterogeneous databases.

Background Art

Information sources such as databases, spreadsheets, tables, and the like are well known in the art. As used herein, "information sources" is intended to broadly refer to data sets that allow some form of querying, either directly or indirectly, with an example of indirect querying as through a suitable "wrapper" layer that functions as a converter or interpreter. Information source examples include, but are not limited to, tables, databases, spreadsheets, web pages with or without forms, flat files, software with API's (application program interfaces), and the like.

When interpreting such a data set, there are different levels of understanding involved. For example, any data can be viewed at the bitstream level (zeros and ones) and the character level (e.g., ASCII or Unicode). Apart from this very low, close to physical level representation of information, higher level structures like records, sets, lists, trees, graphs, etc. are employed to provide better abstractions and handles for data and information manipulation. For example, a relational database hides its physical data organization from the user and only exposes a logical view of the modeled "mini world." (every database can be seen as a representation of some aspects of the world, hence the term "mini world").

This logical view is captured in the *relational database schema*. This comprises, for each relation, the relation name (= table), and the names and data types of the relation's attributes (= table columns). In addition to this already high-level logical view, there is a higher "conceptual level view" on the database which is often not made available to the user, either because there does not exist a formal (machine-readable) representation of that conceptual level view, or, even if it exists, e.g., in the form of an entity relationship ("ER") or uniform modeling language ("UML") diagram, this representation may not be linked to the database query mechanism in a systematic way.

This highest conceptual level representation of databases may be characterized in a conceptual model, often in a language such as ER or UML. A conceptual model represents knowledge that is not discernable from the face of the data source. For example, a data source's conceptual model can represent implicit "domain rules" (or "domain semantics") that capture additional aspects of the source's modeled mini world

By way of a simple example, assume a car manufacturer X is interested in answering a question of the form "which parts of the 1998 'Hector SUV' were purchased or serviced most between January 1, 1999 and December 31, 2000". A prior art database approach can answer such questions based on tables of the form:

Table Sold Parts

Transaction ID	Customer ID	Date	Vehicle ID	Model	Year	Part No	Qty	Price
----------------	-------------	------	------------	-------	------	---------	-----	-------

Table Services

Transaction ID	Customer ID	Date	Vehicle ID	Service ID	Service Description	Price
----------------	-------------	------	------------	------------	---------------------	-------

Next, consider that car company X has a database at its production site that keeps track of the jobs performed by different machines in the shop floor. In addition to the job performed at each machine, the database keeps track of a

"machine's health" by recording its service dates and errors produced by it. This information is maintained using the following simplified schema:

Table Jobs

Machine ID	Job ID	Date	Time	Part No	Error	Comments
---------------	-----------	------	------	------------	-------	----------

5

Table Machines

Machine ID	Machine Type	Operation	Last Service	Next Service	Maximum Error	Comments
---------------	-----------------	-----------	-----------------	-----------------	------------------	----------

With this schema a query such as "which parts were produced with machines whose cumulative error exceeded the maximum error before its service date, such that the date of the part production by the machine is after the date the machine exceeded the maximum error?" can be processed. The query result will represent "defective parts" that might have been produced by "defective machines".

However, these two data sources, even if integrated, may not be useful in processing other queries. For example, consider queries directed to determining whether the defective parts produced by defective machines have any relationship with the parts that are purchased or serviced most in the parts shops. At a first glance it may appear that the two information sources could be joined with reference to their part numbers to process such queries. Such an approach, however, will produce only incomplete results. In particular, only those parts that were both defective and were serviced or purchased would be identified. Intuitively, this result is incomplete in that a specific defective part, say in the transmission, may not need any service at all but may instead cause other parts it interacts with to require service/replacement.

These example car-related data source and queries are fairly simple. Data sources may be much more complex, depending on the complexity of the mini-world they represent. Also, sophisticated data sources often have

sophisticated query capabilities. Such sophisticated data sources may be found in the area of biological research, for example, where a genomic database may have the ability to search large amounts of genomic data to report similar gene sequences using complex and specialized string matching algorithms. As
5 another example, macromolecular databases compare the 3D structure of molecules to determine their possible structural relationships.

As databases and other data sources have become more powerful and widely used, users are often faced with the task of obtaining information from a plurality of sources. Once again referring to the art of biological research by
10 way of example, a biologist may assess different animal models to study different aspects of the same biological function. Thus, a biologist may wish to integrate, for instance, information from a first database regarding the brain from a rodent, from a second database regarding the brain from a primate, and
15 from a third database regarding portions of primate and rodent brains that deal with vision. All three of these databases may have been created at different times by different researchers using different models (or they may in fact come from one common database/design process as in the car manufacturer example above). In particular, each individual database may have different semantics resulting in different structures, and have different query capabilities. As a
20 result, there are numerous difficulties associated with attempting to universally query the databases.

Solutions to these difficulties have been proposed. For example, so called "mediator systems" have been offered to integrate data from different data sources. FIG. 1 is a schematic generally illustrating a prior art mediator
25 system architecture. The mediator generally accesses data from the various databases by means of "wrappers" which sit "on top" of the sources and create a uniform access mechanism to them. The wrappers generally export data from the database in a common, often so-called "semistructured" language, so that any data (whether it is very structured like data from databases, or whether it is
30 less structured such as certain HTML documents) from the various sources can

be presented to the mediator in a uniform data language. A prominent semistructured data language is the Extensible Markup Language XML.

The user query is generally translated by the mediator into an XML query when issuing requests downwards, and XML result elements when
5 sending back data upwards. The wrapper layer at each database translates the incoming XML query into a language native to the database query capabilities. Results obtained from each database may then be conveyed to the mediator in XML and presented through the user interface to the user.

For more information regarding mediator systems, reference may be
10 made to D. Florescu, L. Rashid, and P. Valduriez, "A Methodology for Query Reformulation in CIS Using Semantic Knowledge", Intl. Journal of Cooperative Information Systems, vol.5, no. 4, pp. 431-468, 1996, World Scientific Company; H. Garcia-Molina, Y. Papakonstantinou, D. Quass, A. Rajaraman, Y. Sagiv, J. Ullman, V. Vassalos, and J. Widom, "The TSIMMIS
15 Approach to Mediation: Data Models and Languages", Journal of Intelligent Information Systems, vol. 8, no. 2, 1997, Kluwer Academic Publishers; V. Kashyap and A. Sheth, "Semantic and Schematic Similarities between Database Objects: A Context-based Approach", VLDB Journal, vol. 5, no. 4, pp. 276-304, 1996, VLDB Endowment, Saratoga, California; and Springer-
20 Verlag; L. Haas and D. Kossman and E. Wimmers and J. Yang, "Optimizing Queries across Diverse Data Sources", In Proc. International Conference on Very large Databases, Athens, Greece, pp. 276-285, 1997, VLDB Endowment, Saratoga, California.

Such prior art systems have proven useful when combining different
25 data sources whose relevant (for the integration) classes, or "inter-source couplings" ("ISC"s) are more or less evident from the native source schema. For example, common, similar, or very closely related attribute names may indicate "joinable" columns – a very common ISC (e.g., the above relational database schemas, where "part number" may provide a relevant ISC). By way
30 of further example, mediator systems as known have proven useful for applications such as comparison shopping for a particular appliance model on

the world wide web where different vendors may use databases that have different structures, yet whose semantics make it fairly simple to integrate the sources to process a query using simple ISC's. In this example, it is straightforward to search for the appliance model number in the different
5 databases and combine data from several databases into a single set.

For more complex integrations, however, prior art systems have achieved only limited usefulness. For example, different data sources may be difficult or impossible to integrate with systems and methods of the prior art if the data sources have relations that are not "apparent" to the mediator tool and
10 that have not been encoded in them.

For these and other reasons, unresolved needs in the art exist.

Disclosure of the Invention

The present invention is directed to integration methods and computer
15 program products for integrating data from a plurality of data sources. In particular, the present invention is directed to methods and computer programs that integrate semantic information from the plurality of data sources into an integrated semantic model. A method of the present invention has general steps of obtaining semantic information from each of a plurality of different
20 data sources and creating a conceptual model of the data source that expresses the semantic information. An additional step comprises accessing one or more secondary knowledge sources having knowledge that relates the data from different data sources. An integrated semantic model of the plurality of data sources is then created in an additional method step using the conceptual
25 models of each of the plurality of data sources and the secondary knowledge source. Preferably, the integrated semantic model comprises logic statements that relate classes from different of the data sources.

It will be appreciated that embodiments of the present invention lend themselves well to practice in the form of computer program products.
30 Accordingly, embodiments of the invention comprise computer programs

having computer readable instructions that when executed cause a computer to perform steps of methods of the invention.

The present invention thereby solves many of the heretofore unresolved problems in the art in an elegant and novel manner. By way of example, in the context of the car-company examples presented above, queries such as “which defective parts produced by defective machines have any relationship with the parts that are purchased or serviced most in the parts shops?” were difficult to evaluate with methods of the prior art. To accurately process such a query, it is necessary to have, in addition to the given information source tables, secondary knowledge that bridges the gap between the sources. In particular, relevant secondary knowledge in this example may comprise a set of auxiliary facts and rules about how one part of a car relates to and can impact another part. These secondary knowledge facts and rules can be represented in a formal, machine-readable way, with a preferred example comprising a labeled graph-structure representing logic relations.

The above brief description sets forth broadly some of the features and advantages of the present disclosure so that the detailed description that follows may be better understood, and so that the present contributions to the art may be better appreciated. There are, of course, additional features of the disclosure that will be described hereinafter that form the subject matter of the claims appended hereto. In this respect, before explaining embodiments of the disclosure in detail, it is to be understood that the disclosure is not limited in its application to the details of the construction set forth in the following description or illustrated in the drawings. The present invention may provide additional embodiments, as will be appreciated by those knowledgeable in the art. Also, it is to be understood that the phraseology and terminology employed herein are for description and not limitation.

Brief Description of the Figures

FIG. 1 is a schematic diagram of a mediator system of the prior art.

FIG. 2 is a schematic diagram illustrating in general the architecture of a program product embodiment of the invention.

FIG. 3 is a schematic diagram illustrating the architecture of a second program product embodiment of the invention.

5 FIG. 4 is a graphic representation of secondary knowledge in the form of a domain map useful in an example application of an invention embodiment.

FIG. 5 is a schematic illustrating an example invention embodiment refined secondary knowledge domain map.

10 FIG. 6 is an example input screen from a computer program invention embodiment.

FIG. 7 is an example schematic illustrating an invention embodiment aspect related to presenting output data in context.

FIG. 8 is an example numerical data output screen from a computer program product embodiment.

15 FIG. 9 is a schematic illustrating an example invention embodiment architecture.

Best Mode of Carrying Out the Invention

20 It will be appreciated that the present invention is well suited for practice in the form of computer program products as well as methods. Accordingly, invention embodiments comprise methods as well as computer programs. Indeed, it will be understood by those knowledgeable in the art that method embodiments may be closely related to computer program
25 embodiments in that methods carried out by a computer may comprise computer programs. Therefore, while discussion herein may make reference to one of either a method embodiment or a computer program embodiment, it will be understood that the discussion may be likewise applicable to the other of a method or computer program embodiment.

30 In the case of computer program product embodiments, the computer programs comprise computer executable instructions stored in a computer readable medium that when executed cause a computer to execute steps. Those

knowledgeable in the art will appreciate that a computer readable medium may comprise any of a wide variety of mediums including, but not limited to, magnetic and optical media such as disks and the like, electrical conductor devices, circuitry and micro-circuitry, and the like. Also, it will be appreciated
5 that the term "computer" as used herein is intended to broadly refer to any processor-based device capable of processing computer readable instructions, and by way of example includes, but is not limited to, desktop processor devices, portable processor devices, network based processor devices, and the like.

10 FIG. 2 is a schematic useful in describing in general an embodiment of a method and program product of the invention. In general, a method embodiment of the invention comprises obtaining semantic information from each of a plurality of different data sources 100-104, and creating
15 corresponding conceptual models (CM) 106-110 that model the semantic information of the respective data sources. As used herein, the term "semantic information" as used with data sources is intended to broadly refer to the characterization of constraints that hold for subsets of data and of relationships that hold between data. For example, "semantic information" as used herein is
20 intended to refer to groupings that may be present within data elements, relative numerosity of different categories of data items, rules that a set of data elements must always abide by, or the like. By way of further definition, "semantic information" as used herein may express the properties of data that have not been explicitly encoded in the alphanumeric representation of the data or in the syntactic structure that holds together different data elements.

25 Those knowledgeable in the art will appreciate that semantic information may be obtained in any of a variety of particular steps. For example, for a digital source such as a relational database, semantic information such as constraints and uniqueness of values may be obtained from the "system catalog". Also, it may be possible to infer semantic information by
30 inspecting the structural relationship among data elements. For example, it may be possible to partly infer the semantic information for a data source that

has an additional "document model" enumerating the "tags" into which the document producer has categorized the data. By way of additional example, semantic information may be obtained through personal inspection or observation of a data source. That is, often semantic information is apparent to
5 a learned data source reader or observer through personal knowledge of the domain that gives rise to the data. Other particular methods of "obtaining" semantic information will be apparent to those knowledgeable in the art.

Also, as will be appreciated by those knowledgeable in the art, the term "conceptual model" as used herein is intended to broadly refer to a model for
10 expressing conceptual level semantic information of a data source that is not available on the face of the data source. By way of particular example, a conceptual model may be expressed by an (E)ER diagram, a UML class diagram, or by a Resource Description Framework (RDF).

Method and program product embodiments of the invention further
15 comprise accessing a secondary knowledge source 112 that has information relating at least some of the data sources 100-104 to one another. The term "secondary knowledge" as used herein is intended to refer to information such as facts, constraints, rules, or the like that relate data from the data sources 100-104 to one another. The secondary knowledge source 112 may be thought of
20 as a semantic "roadmap" that connects the seemingly unrelated data of the various data sources 100-104 to one another. A secondary knowledge source 112, by way of example, may comprise expert knowledge or a digital data file that can relate data from one or more of the data sources 100-104 to data from other of the data sources 100-104 by logical "is a", "has a", "part of",
25 "connected to", or the like relations (e.g., B "is a" C, B "has a" D, D "is a part of" E, etc.).

Several of these logical relations may also express their properties through rules. For example, a logic rule may state that for a specific domain, a "has a" or a "part-of" relation is irreflexive, anti-symmetric, and transitive.
30 Similarly, the secondary knowledge source may also specify rules that state joint properties of multiple relations. For example, a rule demonstrating the

interplay of "is a" and "has a" relations may be: "If A is a B, and B has a C, then infer that A has a C". Similarly, there may be several different types of "has a" relationships that may or may not be interrelated as specified by rules. Using the secondary knowledge source 112 and the conceptual models 106-
5 110, a mediator 114 creates an integrated semantic model 116 that provides a single, "integrated semantic model" of the data from the different data sources to a user.

The term "semantic model" is intended to broadly refer to a conceptual level model that expresses the semantics of the underlying object being
10 modeled. Further, the term "integrated semantic model" as used herein is intended to broadly refer to a conceptual level model that relates data from different sources. An integrated semantic model may also be referred to herein as a domain map. In either case, the integrated semantic model Further, application of a query against the integrated semantic model or the domain map
15 may result in identification of data relevant to the query. This set of relevant data may be referred to herein as an integrated view or an integrated view definition.

With the relatively straightforward discussion of the invention embodiments as illustrated by FIG. 2 complete, the more detailed invention
20 embodiment of FIG. 3 may be discussed. The schematic of FIG. 3 illustrates the architecture of a preferred computer program product of the present invention. In this architecture, a plurality of different databases 200-204 are to be integrated for processing one or more queries. To accomplish this integration, the program product obtains data from the databases in a
25 semistructured representation. As used herein, the term "semistructured" is intended to refer to a unified data representation comprising instance (i.e., data-level) and schema-level information (i.e., schema information such as "attribute name is directly represented in the data").

In semistructured data, for instance, the data may be represented by a
30 labeled, ordered tree or by a labeled, directed graph where nodes of the structure represent the actual data, and the labels of the edges represent the

"descriptor" of the data at the end of the edge. By way of example, an XML document can be considered to be "semistructured" data. A relational table may also be captured in a semistructured representation. For example, a portion of the data in the Table named Jobs from the earlier example is represented as:

Machine ID	Job ID	Date	Time	Part No	Error	Comments
1234	109	1/2/98	13:42	82-53	0.02	OK

In a semistructured representation using XML the same information can appear as:

```

10      <jobs>
        <job job_ID="109">
          <machine_ID>1234</>
          <date>
15              <month>1</>
              <day>2</>
              <year>98</>
          </date>
          <time>13:42</>
          <part_no>82-53</>
          <error>0.02</>
          <comments>OK</>
20      </job>
      </jobs>

```

Since the flexibility of semistructured representation can encode many different data models, it is preferred that the data sources export their information to the integration unit (i.e., the mediator 217) in XML. The task of converting the data from the data source's native representation to XML is performed by software modules schematically described as "XML wrappers" 206-210 in FIG. 3.

In addition to wrapping data in the semistructured XML form, a computer program embodiment also obtains semantic information that may comprise class structure (e.g., a class hierarchy or a class lattice), relationship schema, semantic rules, and the like from the databases to create a conceptual model for each database. Indeed, an important aspect of the present invention is that conceptual models themselves can be exported to the mediator 217.

Conceptual models are conveyed using software modules 212-216 referred to as the conceptual model wrappers (CMWs). The task of the CMW's 212-216 is to convert the conceptual model from the data source's formalism (e.g., EER or UML) to a uniform and consistent formalism.

5 The semantic information that comprises the conceptual model is conveyed by registering CMW's 212-216 with a mediator 217. Registration requires that each CMW 212-216 convey semantic information such as class schema, relationship schema, semantics rules, and the like to an interface layer at the mediator in a language that the mediator 217 can interpret and operate
10 with. At this mediator interface layer, the generic conceptual models ("GCM's") 218-222 will be created in the consistent and uniform language.

 Preferably, the GCM's 218-222 are object oriented models that comprise first order logic declarations, and most preferably have the form of frame logic ("F-logic"). Those knowledgeable in the art will recognize that
15 many suitable equivalents for practice with the invention are available in addition to F-logic, including but not limited to Datalog and Prolog. Those knowledgeable in the art will also appreciate that representing the GCM in F-logic imparts well founded semantics to the logic rules. The principles of well-founded semantics can be appreciated by any one knowledgeable in the art of
20 logic programming. The use of first order logic naturally enables the GCM to express functionalities such as cardinality constraints, range constraints, inclusion dependencies, and the like; as well as useful inductive constraints such as transitivity of the class hierarchy, and the like, when using well-founded or similar fixpoint based semantics.

25 The logic statements of the GCM's 218-222 may express declarations of at least instance, subclass, method, and relation. An instance (X, C) specifies that the object named X is an instance of class C. With reference to the auto company X example presented above, "part 389" is an instance of the class "Parts". A subclass (C₁, C₂) specifies that C₁ is a subclass of C₂ and that
30 instances of C₁ are therefore instances of C₂. Again by way of example, "Moving Parts" are a subclass of "Parts", and instances of the subclass

"Moving Parts" such as "Piston2" are therefore instances of the class "Parts". The method (C, M, C_M) specifies that the method M is applicable to objects in C to yield objects in C_M . The method (fuel_pump, supplies, fuel_lines), for instance, specifies that for the class fuel_pump, the method "supplies" produces
 5 result of the type "fuel_lines". The relation $(R, A_1/C_1, \dots, A_n/C_n)$ specifies an n -ary relationship between objects of classes C_1, \dots, C_n , with A_n corresponding to attributes or association roles.

In addition to these logic declarations that logically relate data from different databases, GCM's 218-222 further preferably comprise additional
 10 logic rules referred to herein as integrity constraints. The integrity constraints serve as logic "checks" to insure that information and relations deduced through the logic declarations of GCM's 218-222 are consistent. For example, assume a logic constraint ϕ should hold for all instances of a CM. By expressing it as the denial $\psi := \neg \phi$ (i.e., a named error has occurred if the
 15 negation ψ of ϕ holds). Hence if ψ can be derived for a particular CM instance, then the CM instance violates the ϕ and an inconsistency has been detected. Inconsistencies are tabulated as they are encountered to monitor consistency and indicate and explain inconsistency.

Query capabilities may likewise be conveyed from CMW's 212-216 to
 20 the mediator 217 for inclusion in the GCM's 218-222. A data source's query capability may be illustrated with reference to an example form used on a web page. Assume that a user can query the table named "Services". While the table has the structure:

Transaction ID	Customer ID	Date	Vehicle ID	Service ID	Service Description	Price
----------------	-------------	------	------------	------------	---------------------	-------

25 the user can only access it through a form that shows the fields "Customer ID", "Date", "Vehicle ID" and "Service ID". By filling out these fields, the user can retrieve a corresponding record. In this example, the table query capabilities do not allow for the user to formulate a query such as "Find all records where the price is less than \$50". However, if the source had allowed
 30 the user to ask an SQL query, this will be a very simple query to formulate.

Thus, two sources with identical logical and conceptual models may have different query capabilities.

When the CMWs 212-216 convey their conceptual-level semantic information (e.g., conceptual models expressed in languages such as EER, RDF, UML, or the like, and additional rules describing the semantics of source data) to the mediator 217, they preferably do so using an XML representation, for example, a DTD for encoding EER or UML (e.g., XMI is such a language for UML). If the mediator 217 already "understands" XMI (i.e., capable of recognizing and mapping XMI constructs into equivalent GCM construct), then creation of a GCM 218-222 is straightforward. If the CMW 212-216 exports its semantic information in a language new to the mediator 217 and thereby not "known" (e.g., a new language "NCMF" that may be an alternative XML representation for UML models), then a "CM plug-in" may be provided or performed.

The ability to use "plug-ins" for registering otherwise unknown CM languages with the mediator advantageously provides for excellent extensibility of the program product architecture. Additional data sources may be readily added to the dialog regardless of their CM formalism, so long as the formalism may be converted to the GCM logic language. Once a plug-in has been provided, the mediator will be able to recognize and convert future instances of the CM representation. The preferred GCM object oriented and F-logic representation allow for the GCM's 218-222 to formalize virtually any CM formalism used. CM representations such as UML, RDF (resource description framework) or EER (extended entity relationship), for example, may be readily converted to uniform object oriented F-logic at the GCM level.

As an example of the plug-in aspect of the invention, assume the mediator 217 uses frame logic ("F-logic") as a GCM representation. In order to register with the mediator 217, a new data source using an unknown language NCMF to encode its UML conceptual model (including semantic domain constraints) has to be translated to equivalent F-logic facts and rules. Hence an NCMF-to-F-logic plug-in is provided. Formally, a plug-in is a

mapping "f" from the source's CM language (here: UML in NCMF) to an equivalent representation in the GCM language (here: F-logic) of the mediator. For example, a mapping:

$f: \text{UML (here: in NCMF)} \rightarrow \text{GCM (here: F-logic)}$

5 that translates any given UML class diagram (in NCMF encoding) into an equivalent F-logic representation would be such a plug-in. In this manner, the UML semantics of a source S is captured in NCMF and after applying f it is represented in the form of executable F-logic statements (rules) in a GCM. Hence the data source's ("S") semantic information (as captured in S.NCMF)
10 can be employed for creation of a GCM for use by the mediator. The data source S may then have a GCM expressing an equivalent F-logic ("FL") representation, which may be summarized as $S.FL := f(S.NCMF)$.

Should the mediator need to issue a query Q against the GCM of the data source S (i.e., S.FL), the query Q may be translated into an equivalent
15 query Q' against S.NCMF. Consider:

$Q' := Q \circ f$

i.e., the composition comprising an application of f (going from NCMF to GCM),

followed by an application of Q (going from GCM to the result schema).
20 Clearly this Q' achieves the desired effect. Thus, a possible, albeit probably inefficient, way of evaluating Q' is to compute the complete image of NCMF under f and then to apply Q to it. Note that this amounts to materializing, at the mediator, the semantics of S in the GCM language. An often preferable solution is to not execute f and Q in sequence but to compute the functional
25 composition $(Q \circ f)$ a priori. If such a composition is used which does not need to materialize the f image, then this often results in a more efficient Q'.

In the computer program embodiment architecture as represented by the schematic of FIG. 3, the mediator 217 contains a domain map 224 representing the integrated semantic model created through consultation with a secondary
30 knowledge source that has information relating the individual GCM's 218-222. That is, the secondary knowledge source provides the "semantic road map" or

“knowledge glue” that bridges the gap between the GCM’s 218-222 that may be formalized in the domain map 224 to provides facts, rules, constraints, and the like relating the various GCM’s to one another. Embodiments of integrated semantic models such as the domain map 224 thereby formalize expert
5 knowledge that is needed to mediate across multiple world scenarios as represented by different information sources. Preferably, the domain map 224 comprises first order logic statements in the same language with those of the GCM’s 218-222, and therefore preferably comprise F-logic declarations.

By way of example, the logic statements of the domain map 224 may
10 describe relationships between members of the different databases such as instance, subclass, method, relation, and the like. By way of example, relations between classes such as “is a”, “has a”, “part of”, “process of”, “before”, “happens after”, “located in”, “connected to”, “less than”, “greater than”, and the like may be described. Preferably, each such logic statement has a set of
15 logic rules that describes the properties of the statement. For example, statements may be described as transitive, reflexive, symmetric, or the like. Logic constraints, ranges, and the like may also be expressed.

By way of further example, an integrated semantic model in the form of a domain map for the car company X example presented above may specify the
20 domain knowledge that “pouring excessive fuel through the fuel injector may reduce the fuel’s ability to ignite properly”. A different rule may state that “the nozzle diameter of the fuel injector, the pressure of the fuel from the fuel line, and the pulse-width of the electromagnetic energizer of the fuel injector regulate the amount of fuel delivered to the engine”. The domain knowledge
25 may also represent “connectivity” information such as “the electromagnetic energizer of the fuel injection system is connected to (and controlled by) the engine control unit, which is a computer, and that every engine has an engine control unit and an air pressure sensor”.

In an invention architecture embodiment, integrated semantic models as
30 represented by domain maps (“DM”s) may comprise special conceptual models whose class names denote *concepts*. Concepts provide the “semantic anchor

points" from which sources can "hang off" their data. Concepts can be linked via binary relations called *roles*. Intuitively, a labeled edge $C \rightarrow^r D$ of a DM means that if $c \in C$ then there is some $d \in D$ such that $r(c,d)$ holds: e.g., $\text{dendrite} \rightarrow^{\text{has}} \text{branch}$ means that every dendrite must have some branch.

5 More particularly, a DM embodiment may comprise a finite set comprising (i) description logic *facts*, and (ii) *logic rules*, both involving finite sets C (*concepts*) and R (*roles*). Facts are visualized as *edge-labeled digraphs* where nodes correspond to concepts and edge labels to roles. In this DM embodiment, edges expressing necessary conditions for inclusion of C in D (is
10 a) and existence of r -related objects in D (*ex*) are most common. If all r -related objects of any $c \in C$ have to be in D , (all) is used. Nodes AND and OR correspond to *conjunctions* and *disjunctions* of the union of all outgoing edges, respectively. When unique, AND nodes are omitted and outgoing arcs directly attached to the concept are defined. Necessary and sufficient conditions of C
15 are expressed using (*eqv*).

Description logic is a subset of first-order logic ("FO"). There are different ways in which we can "execute" the axiom for an edge at the mediator, for example as an *integrity constraint* or as an *assertion*: e.g., translating $\text{FO}(\text{ex})$ as an integrity constraint in F-logic (FL) tests whether the
20 mediator's object base contains for each $X : C$ a corresponding $Y : D$; otherwise a violation is reported. Such an integrity constraint is useful when the mediated object base is required to be data-complete wrt. $C \rightarrow^r D$.

The other, more frequent case is to view $\text{FO}(\text{ex})$ as an assertion that in the real world (but not necessarily in the object base) the corresponding target
25 object y exists. The following F-logic assertion creates a virtual placeholder object $\text{fC};r,D(X)$ whenever the object base does not contain y :

$$Y : D, r(X, Y) :- X : C, \text{not}(Z : D, r(X, Z)), Y = \text{fC};r,D(X).$$

If necessary, a DM can use the full expressive power of the underlying GCM. However, expressiveness is paid for with complexity and the
30 requirement

(EXPR) above can make *reasoning about concepts* undecidable: Subsumption and satisfiability are undecidable for unrestricted GCM domain maps. Here, *subsumption* means to decide whether membership in a concept class C implies membership in another class D, for all logic interpretations (i.e., instances of the DM) I that satisfy a given domain map DM. Satisfiability is the question whether such an I exists. Experience suggests that in a typical mediator embodiment of the invention, *reasoning about the DM* may be required only to a limited extent. Instead, a specific DM is given to navigate the "multiple worlds" and to define and *execute* integrated views.

10 With reference drawn to the embodiment architecture of FIG. 3 once again, to process a query input by a user through a user interface 226, a program product embodiment of the invention will use a mediator engine 228 that preferably comprises an F-logic query engine. Examples of available F-logic engines that can evaluate F-logic expressions are FLORA and FLORID.
15 The mediator engine 228 interfaces with the domain map 224 to create the integrated view definition 230 that represents the populated integrated semantic model from the three databases 200-204 responsive to the input query. For example, in the car company X example case, if the machine producing the nozzle of the fuel turned out to be producing larger-than-permitted nozzle
20 diameter, and the parts most serviced turned out to be the fuel-injection unit, then the original query can be answered faithfully by using the domain knowledge illustrated above.

 In processing a user query through the architecture of FIG. 3, some query and data conversion and filtering may occur. The query as input from
25 the user, for instance, may be converted to a language consistent with the domain map 224 and the GCM's 218-222 (e.g., F-logic). The query, or portions of the query, may likewise be converted from the GCM language (e.g., F-logic) to the language native to the particular CMW 212-216 (e.g., UML). The query may likewise be partitioned into sub-queries appropriate for each
30 CMW. At the CMW 212-216 level, the thus converted query may then be

specified in a query language native to the database. The retrieved results may likewise be ultimately converted to a language consistent with the GCM.

Embodiments of the program product and method of the invention also allow for results obtained from query processing to be viewed in context. That is, results may be presented along with a summary of the logical relations that were pursued to obtain those results. By way of example, assume that a query is answered by applying the query to an integrated semantic model and accessing multiple data sources as guided by the integrated model's relation information. In such a case the mediator may create a graph or other representation that summarizes all the classes and logic relations used during the query processing. This graph will then be made available to the user in addition to data responsive to the query, so that the user can see in what context the data was determined to be responsive. Such a contextual representation of data may be referred to as an "integrated view".

For example, the result graph for the illustrative car-query "which defective parts produced by defective machines have any relationship with the parts that are purchased or serviced most in the parts shops?" would show how the defective part (e.g., the nozzle), the corresponding defective machine, an influence diagram relating the nozzle to the fuel injection unit and the rate of service of the fuel injection unit are correlated.

In order to further illustrate the best mode of practice of the present invention, examples of practice of embodiments of the invention will be presented and discussed herein.

Example 1

For purposes of Example 1, assume that three object oriented databases are to be integrated. The first database concerns employees in a municipal streets and sanitation department, and lists department employees by name, job title, and date of hire:

Streets and Sanitation		
Name:	Job Title:	Date of Hire:
....
Davis, Jim	Driver	2-13-90
Enright, Mike	Helper	6-15-97
Franklin, Ed	Driver	6-15-97
....

A second database concerns the police department for the municipality, and lists officers by name, precinct, and birth date:

Police Department		
Officer:	Precinct:	Birth Date:
....
Alex, Jim	1	2-13-70
Baker, Mike	1	6-15-67
Conroy, Ed	2	6-15-57
....

- 5 The third database concerns the water department, and lists employees by name, job title, and birth date:

Water Department		
Employee:	Job Title:	Birth Date:
....
Crabb, Jim	Supervisor	1-3-70
Douglas, Ted	Laborer	5-1-69
Zlotkowski, Ed	Driver	7-6-65
....

Assume that these three databases are to be integrated to process the query "How many city workers have a valid driver's license?"

- Semantics information obtained from the first database may comprise
- 10 the knowledge that Streets and Sanitation workers are listed by name, title, and date of hire, with each employee record comprising data members in each of these classes, and with a conceptual model showing the same. Similarly, the conceptual model for the police department database may show that data classes exist for "officers", "precinct", and "date of birth". Finally, the water
- 15 department database conceptual model may show that it contains data classes for "employee", "job title", and "birth date".

To create an integrated semantic model of these three databases sufficient to process the query: "how many city employees have a valid driver's license?", secondary knowledge is required to logically connect the classes of the three databases, as well as to provide other data not available from the databases. By way of example, assume expert knowledge is available to show that 1) all police officers have a driver's license, 2) all drivers in the streets and sanitation department have a valid driver's license, 3) all drivers in the water department have a valid driver's license, and 4) all members of the police department, water department, and streets and sanitation department are city employees.

With this secondary knowledge, an integrated view of the three databases can be created for processing the query. The view can illustrate the logic inferences made to process the query, so that the user may view resultant data in the context from which it comes. Now, the original query, "How many city workers have a valid driver's license?" will be transformed into the following query plan:

At the mediator: Unfold the view "city worker" as the union of the three data sources.

From "Streets and Sanitation": count all employees whose job title = "Driver"; assign to variable A at mediator.

From "Police Department": count all employees; assign to variable B at mediator.

From "Water Department": count all employees whose job title = "Driver"; assign to variable C at mediator.

At the mediator: compute result = A+B+C

Example 2 (Two Neuroscience Worlds)

A second example invention embodiment comes from the biological arts, and shows an invention embodiment in a somewhat complex real world application.

Consider two Neuroscience laboratories that perform experiments on two different brain regions. The first laboratory, SYNAPSE, studies dendritic spines of pyramidal cells in the hippocampus. The primary schema elements are thus the anatomical entities that are reconstructed from 3-dimensional serial-sections. For each entity (e.g., spines, dendrites), researchers make a number of measurements, and study how these measurements change across age and species under several experimental conditions.

In contrast, the NCMIR laboratory studies the Purkinje Cells of the cerebellum, inspecting the branching patterns from the dendrites of filled neurons, and localization of various proteins in neuron compartments. The semantic information from the data source indicates that schema used consists of a number of measurements of the dendrite branches (e.g., segment diameter) and the amount of different proteins found in each of these subdivisions. Assume that semantic information from each of the databases indicates that each of the two schemas has a class C having a location attribute that can have a value like "Pyramidal Cell dendrite" and "Purkinje Cell", respectively.

How are the schemas of SYNAPSE and NCMIR, as represented in their generic conceptual models, related? Evidently they carry distinctly different information and do not even enter the purview of the schema conflicts usually addressed in database integrations of the prior art. To an expert scientist however, they are related through the following reason: Release of calcium from spiny dendrites occurs as a result of neurotransmission and results in changes in spine morphology (sizes and shapes obtained from SYNAPSE). Propagation of calcium signals throughout a neuron depends upon the morphology of the dendrites, the distribution of calcium stores in a neuron and the distribution of calcium binding proteins, whose subcellular distribution for Purkinje cells are measured by NCMIR.

Thus, a researcher who wanted to model the effects of neurotransmission in hippocampal spines could get structural information on hippocampal spines from SYNAPSE and information about the types of calcium binding proteins found in spines from NCMIR. Note that in order to

connect the two sources, we need, independent of the observed experimental source data available from the data sources, secondary knowledge like the following:

- 5 • *Purkinje cells and Pyramidal cells have dendrites that have higher-order branches that contain spines .*
- *Dendritic spines are ion (calcium) regulating components.*
- *Spines have ion binding proteins.*
- 10 • *Neurotransmission involves ionic activity (release). Ion-binding proteins control ion activity (propagation) in a cell. Ion-regulating components of cells affect ionic activity (release).*

This secondary knowledge may be provided, for instance, by an expert scientist, a data file, or the like. To formalize this secondary knowledge we can use *description logic* (DL) statements such as:

- Neuron $\sqsubseteq \exists$ has.Compartment
- 15 Axon, Dendrite, Soma \sqsubseteq Compartment
- Spiny Neuron \equiv Neuron $\cap \exists$ has.Spine
- Purkinje Cell, Pyramidal Cell \sqsubseteq Spiny Neuron
- Dendrite $\sqsubseteq \exists$ has.Branch
- Shaft \sqsubseteq Branch $\cap \exists$ has.Spine
- 20 Spine $\sqsubseteq \exists$ contains.Ion_Binding_Protein
- Spine \sqsubseteq Ion_Regulating_Component
- Ion Activity $\sqsubseteq \exists$ subprocess_of.Neurotransmission
- Ion_Binding_Protein \sqsubseteq Protein $\cap \exists$ controls.Ion_Activity
- Ion_Regulating_Component $\equiv \exists$ regulates.Ion_Activity
- 25 Together with additional inference rules (e.g., for transitivity of *has*), DL axioms like these formally capture the secondary knowledge. FIG. 4 shows a graphic portrayal of these DL formulas in the form of a domain map, a kind of semantic net that may be used for defining and executing integrated queries at the mediator.
- 30 This example from a real-world biological science integration scenario illustrates a novel aspect of the present invention over the prior art. In

particular, in the present invention seemingly unconnected schema can be semantically related through secondary knowledge, which in effect situates the data in the scientific context.

5 Example 3 (Adding a New Data Source)

Example 3 is presented to further illustrate the extendability of invention embodiments. As discussed above, a novel advantage of the present invention is that new data sources may be readily added to existing federations of integrated data sources. New sources may be added through registration of
 10 their conceptual models with the mediator. If the mediator does not "recognize" the semantic representation of the new data source, a plug-in converter may be provided for converting to the recognized GCM uniform language.

Thus new data sources may bring their own semantic information including, but not limited to, constraints. After registration of the new data
 15 source, the mediator *refines* the integrated semantic model domain map DM to reflect the new data classes and semantic information.

By way of example, FIG. 5 illustrates a refined DM. FIG. 5 depicts a DM after registering two new concepts MyNeuron and MyDendrite. The
 20 knowledge about these is sent to the mediator using data logic axioms:

$$\text{MyDendrite} \equiv \text{Dendrite} \cap \exists \text{exp.Dopamine_R}$$

$$\text{MyNeuron} \subseteq \text{Medium_Spiny_Neuron}$$

$$\cap \exists \text{proj.Globus pallidus external}$$

$$\cap \forall \text{has.MyDendrite}$$

Thus instances of MyDendrite are exactly those dendrites that express Dopamine R(eceptor), and MyNeuron objects are medium spiny neurons projecting to Globus Pallidus External and *only* have MyDendrites. Assuming
 30 properties are *inherited* along the transitive closure of *isa*, it follows that MyNeuron, like any Medium Spiny Neuron projects to certain structures (OR

in FIG. 5). With the newly registered knowledge, it follows that MyNeuron definitely projects to Globus Palladius External. If we want to specify that it *only* projects to the latter, a *nonmonotonic inheritance*, e.g., using FL with well-founded semantics can be employed.

5

Example 4 (Checking Inductive Properties)

Example 4 is provided to better illustrate an invention embodiment aspect related to checking the consistency and integrity of logic declarations used in GCM's and/or secondary knowledge sources of the invention. This
10 example uses logic language of F-logic ("FL").

The following integrity constraints test whether a binary relation R is a partial order on a class C : rule (1) finds all X in C for which R is not reflexive. Similarly, (2) reports missing transitive edges, and (3) derives node object pairs that violate R 's anti-symmetry on C . Thus, R is a partial order on C iff (1-3)
15 do *not* insert a failure witness into integrity class ic:

(1) $w_{rc}(C, R, X): ic :- X : C, \text{ not } R(X, X).$

(2) $w_{tc}(C, R, X, Z, Y): ic :- X, Y, Z : C,$
 $R(X, Z), R(Z, Y), \text{ not } R(X, Y).$

(3) $w_{as}(C, R, X, Y): ic :- X : C,$
20 $R(X, Y), R(Y, X), X \neq Y.$

If we assign "::" and the meta-class "class" (holding all class names) to the relation variable R and class variable C respectively, the above rules test whether

"::" is indeed a partial order. This example also illustrates the power of schema
25 reasoning in FL.

Example 5 (Cardinality Constraints)

Example 5 is presented to better illustrate logic cardinality constraint aspects of embodiments of the invention. Example 5 uses FL language logic
30 statements, and is set in the data context of the biological arts.

Aggregation and cardinality constraints are ubiquitous in real applications. Consider the GCM declaration $relation(R, A/C_1, B/C_2)$ and assume the CM at hand specifies that the cardinalities of roles A and B satisfy the conditions $card_A(N) := (N=1)$ and $card_B(N) := (N \leq 2)$. Applied to
 5 $has(neuron, axon)$ this says that a neuron can have ≤ 2 axons and an axon is contained in exactly one neuron; more formally:

$w_1(R, VB, N): ic :- N = count \{VA[VB]; R(VA, VB)\}, N = 1.$

$w_{>2}(R, VA, N): ic :- N = count \{VB[VA]; : R(A \quad VA; B \quad VB)\}, N > 2.$

The body of the first rule counts for each value VB of B the number N of values
 10 VA . If $N = 1$ a cardinality violation is detected and the witness w_1 gives the violating triple R, VB, N . The second rule illustrates a different FL syntax for tuple objects and checks $N \leq 2$ for B by grouping on VA .

Example 6

15 Consider a neuroscience laboratory called NCMIR that studies nerve cells of the brain region called cerebellum, inspecting structures within several types of cells, and localization of various proteins in neuron compartments. We can illustrate the schema for this data using tables like:

20 – object declaration for brain section from which the cells were studied
 section (ID, preparedFrom, anatomicLocation, microtome, thickness, fixationMethod);

25 – object declaration for protein localization
 proteinLocalization (ID, proteinName, probeType, probeName, protocolID, probeRefID, section);

– object declaration for protein labeling
 regionOfInterest (regionID, imageID, proteinLocalizationID, points, objectName, area, averageIntensity, varianceIntensity);

30 When a step of an invention embodiment exports this information as a conceptual schema to the mediator, the conceptual wrapper may define a new

exported class called *proteinAmount* instead of the class *regionOfInterest* declared in F-logic syntax as:

5 proteinAmount[proteinName=>string, imageID=> string,
 regionID=>string; objectName=>string; amount=>integer],

which is defined as:

10 proteinAmount[proteinName→P; imageID→I; regionID→R;
 objectName→O; amount→A]:- regionOfInterest.regionID=Y,
 regionOfInterest.imageID=I, regionOfInterest.proteinLocalizationID=
 proteinLocalization.ID, regionOfInterest.objectName=O,
 regionOfInterest.amount=
15 getAmount(regionOfInterest.averageIntensity,
 regionOfInterest.varianceIntensity, regionOfInterest.area, M), M=A.

where *getAmount* is a function that assigns the values in a scale of 0 to 5 based upon some intensity-based computation.

20 Next, consider a second laboratory called, SENSELAB, that determines how one neuron transmits information to another. Let us consider two of their conceptual classes:

25 neuron[name=>string; location=>string; organism=>string; type=string].

and

30 neurotransmission[organism=>string; transmitting neuron=>string,
 transmitting compartment=>string; receiving neuron=>string; receiving
 compartment=>string; neurotransmitter=>string].

Now at the mediator, an integrated semantic model may define a "mediated class" called `protein_distribution` as follows:

5 `D : protein_distribution[protein_name → Y; animal → Z;
distribution_root → P; distribution → D] :- 'NCMIR'. protein.name = Y,
'SENSELAB'.neuron.organism = Z, contains('ANATOM'.
nervous_system.has a star, P), aggregate(Y, 'NCMIR'. protein_amount.
amount, has a star, P, D).`

10 To define this mediated class, a secondary knowledge source called
'ANATOM' has been used which is a taxonomy of the animal kingdom. The
distribution_root P is the name of a brain region like cerebellum such that the
aggregate of the protein Y is computed on starting from this region P. The
predicate `has_a_star` computes the deductive closure over the `has_a` and `isa`
15 relations in 'ANATOM', and is defined with respect to any relation R as
follows:

20 `tc(R)(X,Y) :- R(X,Y).
tc(R)(X,Y) :- tc(R)(X,Z), tc(R)(Z,Y).
dc(R)(X,Y) :- tc(isa)(X,Z), R(Z,Y).
dc(R)(X,Y) :- R(X,Z), tc(isa)(Z,Y).
has a star(X,Y) :- dc(has a)(X,Y).`

25 The function "aggregate" recursively traverses a binary relation R (here:
has a star) starting from node P, and computes the aggregate of the specified
attribute at each level of the relation R.

30 Now using these "integrated views" (in the form of mediated classes
form the integrated semantic model) and the exported classes from the sources,
we can answer the following query:

"What is the distribution of those calcium- binding proteins that are found in neurons that receive signals from parallel fibers in rat brains?"

In terms of the given views, this user query can be written as:

```
5
answer(P, D):- neurotransmission[ organism→'rat'; transmitting
compartment →'parallel.fiber'; receiving neuron→X; receiving
compartment→Y], D : protein distribution[protein name→P; ion_bound
→→ {calcium}; distribution root→_ ].
```

10

This is a typical query of a scientist who studies neurotransmission (and produces the data of SENSELAB above), and needs information gathered by groups that study protein localization (like NCMIR). Note that the user does not specify the distribution root, forcing the mediator to provide a "reasonable" root for the neuron- compartment pairs that satisfy the first condition. The following are the main steps of the query plan executed at the mediator:

1. push selections (rat, parallel fiber) to SENSELAB and get bindings for X and Y
2. using the domain map DM(M), select sources that have data anchored for the neuron/ compartment pairs X, Y from step (1); in our case, only NCMIR is returned
3. push selections given by the X, Y locations to NCMIR, and retrieve only proteins P that are found in X, Y

Now the mediator needs to compute the actual distribution of each protein P from NCMIR. But to do this using the view defined earlier, we first must determine which brain region of the neuron should serve as the root of the

distribution. This is accomplished by computing the least upper bound (lub) of locations in the domain map.

4. with the lub as the root node, compute the view protein distribution at the mediator as described before. Note that this involves a downward closure along the has a star relation. The last two operations filter out a segment in the domain map as the "region of correspondence" between the two information sources, and demonstrate how graph operations on the domain map can be actively used to compute conceptual mappings between sources.

This example demonstrates the actual complexity of the steps involved in correlating conceptually related but schematically distant information sources using the framework of the invention.

- FIGS. 6-8 illustrate graphical interfaces presented through a program embodiment in processing this example query. FIG. 6 illustrates a screen for inputting a query. FIG. 7 illustrates a contextual presentation of results, with classes (ovals) that were related in processing the query connected by relational lines, and results illustrated by the diamond shapes. The embedding of the diamond shape result nodes to data classes illustrates a novel aspect of the invention related to presenting results in a contextual basis. In an invention embodiment, selecting any of these diamonds by, for instance, use of a mouse, can result in display of the numerical data representing the integrated view, with an example display screen shown in FIG. 8 after selection of two diamonds from FIG. 7.

Example 7 (Extensive Mediator Architecture Discussion and Examples)

- Example 7 offers an extensive discussion of the overall scope of an invention embodiment, and it is presented with substantial detail in order to thoroughly illustrate practice of the invention. Example 7 itself contains

several sections and invention embodiment examples that generally are taken form the art of biological science.

A grand goal in many disciplines of biological research is to understand the workings of a biological organ like the brain and how the interplay of different structural, chemical and electrical signals in the biological tissues gives rise to natural and disease processes. To achieve such a goal, however, it is essential to develop an integrated understanding of very different, but conceptually correlated studies and data produced from diverse biological subdisciplines. Most importantly:

- Biologists assess different animal models to study different aspects of the same biological function. Thus, for a given research problem, they may wish to integrate information about the cytoarchitecture of sensory cortex from the somatosensory cortex of the rat, the brain areas involved in vision from the primate, the physiology of receptive fields from the cat, the distribution of key proteins involved from the rat, and the molecular underpinnings of synaptic plasticity from the mouse.
- Biologists study the same biological system from multiple perspectives. For example, in the study of calcium regulation, researcher A may take a physiological approach, using patch electrodes to study calcium currents; researcher B may take an anatomical approach, mapping the distribution of different isoforms of calcium regulatory proteins and the organelles that express them; a biochemist C may study signal transduction cascades and levels of protein activity using Western blots and assay systems, a pharmacologist D may use a panel of channel blockers, agonists or antagonists to study the response in single cells or the whole animal to alterations in calcium regulation.

The goal of this Example 7 is to present invention method and program embodiments that integrate different studies and analyses conducted by biologists performing different experiments, such that the integrated body of information can be queried and navigated across. Once such information is

integrated, the practicing biologist can use the system to discover biologically significant correlations and use the discovery to plan future work in the context of available data.

As is true in many other arts, the integration challenge is that source data cannot be joined using simple term-matching or comparison operators. Even more sophisticated approaches which use ontologies to enumerate joinable terms are often not sufficient. Instead a join should be performed based on whether the objects satisfy some application-specific condition. For complex integration scenarios like a neuroscience application, a more expressive formalism is necessary to specify these "semantic join conditions". In particular, the formalism should have inferencing mechanisms to reason over domain knowledge if necessary.

In summary, Example 7 develops an invention embodiment program and method, referred to within Example 7 as "KIND" (*Knowledge-based Integration of Neuroscience Data*), that represents a substantial improvement over conventional wrapper-mediator methods and programs of the prior art with one or more secondary knowledge sources referred to as *domain knowledge bases* that provide the "semantic glue" between sources through facts and rules from the application domain. Thus our mediator enhances view-based information integration with deductive capabilities. Data manipulation and restructuring operations for integration can be performed not only on the base data from the sources but also on intensional data derivable from the knowledge bases. To this end, we employ the deductive object-oriented language F-logic and demonstrate that it can handle the given complex integration problems.

It should be emphasized that the integration problem addressed by the present invention and the embodiments of Example 7 is different from the problems addressed in prior art efforts that relate to issues like schema integration, resolving conflicts and mismatches (structural, extensional, naming, etc.), global query processing in the presence of local autonomy etc. Those heterogeneities are between different representations of essentially the

same (or very similar) real world entities. In contrast, the present invention deals with sources containing *inherently different* (but related through "expert knowledge") information.

The remainder of Example 7 is organized as follows. Example 7.2 provides a brief introduction to F-logic to clarify the notation and concepts used subsequently. Example 7.3 presents a motivating example that illustrates the nature of the information integration task for the given problem domain. Example 7.4 presents an invention embodiment architecture and explains the role of F-logic in the representation of schema, knowledge and in the inference mechanism. Example 7.5 illustrates a particular instance of the invention embodiment as applied to the Neuroscience art in the form of "INSM" (*Integrated NeuroScience Model*). Finally, Example 7.6 shows in more detail how integrated views are defined and queried in INSM using semantic information.

Example 7.2: F-Logic Summary

Because the invention embodiment of Example 7 is based on F-logic, we briefly introduce the syntax and basic concepts of F-logic. For additional teaching regarding F-logic, reference is made to "Logical Foundations of Object Oriented and Frame Based Languages, M. Kifer, G. Lausen, and J. Wu; Journal of the ACM, 42(4):741-843, July 1995; incorporated herein by reference. While there are other formalisms that could possibly be used, F-logic is preferred for several reasons: F-logic is a declarative language with rich modeling capabilities (class hierarchy, complex objects, inheritance, etc.) and a powerful rule language. It has its roots in AI (frame-based knowledge representation) and deductive object-oriented databases. Apart from "pure" database modeling and querying, it has been applied in several related (but different) areas, including schema transformation, information integration, querying the Web, knowledge representation/reasoning with ontologies, and management of semistructured data.

F-logic extends Datalog and first-order logic (including Skolem functions). In particular, well-known transformations can be used to map arbitrary first-order constraints to equivalent stratified Datalog (and thus F-logic) rules. Finally, F-logic query evaluation engines such as FLORA [FLOa],
 5 FLORID [FLOb], and SILRI are readily available (and continue to be improved).

F-logic Syntax and Object Model

- 10 • *Symbols*: The F-logic alphabet comprises sets F , P , and V of *object constructors* (i.e., function symbols), *predicate symbols*, and *variables*, respectively. Variables are denoted by capitalized symbols (X , $Name$, ...), whereas constants and function symbols (0-ary and n-ary object constructors) are denoted in lowercase (cerebellum, foo(bar,baz), ...) unless quoted ('Cerebellum'). An expression is *ground* if it involves no
 15 variables. In addition to the usual first-order symbols, there are special symbols such as: $]$, $[$, $\{$, $\}$, \rightarrow , $\rightarrow\rightarrow$, \Rightarrow , $\Rightarrow\Rightarrow$, $:$, $::$.
- *Id-Terms/Object-Ids (Oids)*:
 (0) First-order terms over F and V are called *id-terms*, and are used to name objects, methods, and classes. Ground id-terms correspond to
 20 *logical object identifiers (oids)*. In particular, *constants* and *strings* ("cerebellum") are oids; the latter are conceived as *character lists*, i.e., nested ground terms.
- *Atoms*: Let O , M , R_i , X_i , C , D , T be id-terms. In addition to the usual first-order atoms like $p(X_1; \dots; X_n)$, there are the following basic types
 25 of atoms:
 - (1) $O[M \rightarrow R_0]$ (*single-valued meth. app.*)
 - (2) $O[MR_1; \dots, R_{n_g}]$ (*multi-valued meth. app.*)
 - (3) $C[M \Rightarrow T]$ (*single-valued class signature*)
 - (4) $C[M \Rightarrow\Rightarrow T]$. (*multi-valued class signature*)

(1) and (2) are *data atoms* and specify (at the instance level) that the application of *method* M to the object with oid O yields the result object with oid R_i . In (1), M is *single-valued* (or *scalar*), i.e., there is at most one R_0 such that $O[M \rightarrow R_0]$ holds. In contrast, in (2), M is *multi-valued*, so there may be several result objects R_i . For $n = 1$ the braces may be omitted. (3) and (4) denote *signature atoms* and declare that the (single/multi-valued) method M applied to objects of *class* C yields instances of *type* (i.e., class) T . The organization of objects into classes is specified by *isa-atoms*:

- 10 (5) $O : C$ (O is an instance of class C)
 (6) $C :: D$. (C is a subclass of D)

- *Path Expressions*: F-logic supports path expressions to simplify object navigation along single-valued and multi-valued method applications and to avoid explicit join conditions. The following *path expressions* are allowed in place of id-terms:

- 15 (7) $O.M$ (single-valued path expression)
 (8) $O..M$ (multi-valued path expression)

The path expression (7) is *single-valued* and refers to the unique object R_0 for which $O[M \rightarrow R_0]$ holds, whereas (8) is *multi-valued* and refers to each R_i for which $O[M \rightarrow \{R_i\}]$ holds. O and M may be id-terms or path expressions. Although not part of the core syntax, *generalized path expressions* 4 can be defined by means of rules.

- *Parameters*: Methods may be *parameterized*, so

$$M@ (X_1; \dots; X_k)$$

25 is allowed in (1–4) and (7–8).

Example: $o_{rat} [name@ (scientific) \rightarrow "Rattus rattus"]$.

- *Rules*: A rule is of the form: *Head IF Body* where *Head* and *Body* are conjunctions of F-logic atoms (read "if a ground instance satisfies *Body* then also the *Head*"); a program is a set of rules.

30

F-molecules are a concise notation for several atoms specifying properties of the same object: for example, instead of $o_{rat} : \text{taxon} \wedge o_{rat} [\text{name}@\text{(common)} \rightarrow \text{"rat"}] \wedge$

$o_{rat} [\text{order} \rightarrow \text{"Rodentia"}]$ we can simply write

5 $o_{rat} : \text{taxon}[\text{name}@\text{(common)} \rightarrow \text{rat}]; \text{order} \rightarrow \text{"Rodentia"}]$

In F-logic rules “,” is shorthand for “ \wedge ”.

Object Model

An *F-logic database (instance)* is a set of ground F-logic atoms. The
10 basic relations among objects ($\rightarrow, \rightarrow\rightarrow, \Rightarrow, \Rightarrow\Rightarrow$) in this model can be represented as a labeled graph where nodes are oids and where edges are labeled with the corresponding arrow and the method name. From base facts additional facts can be derived by means of rules.

15 **Example 7.2.1 (Fragment of ANATOM)** The following is a fragment of ground F-logic atoms and molecules that make up the anatomical knowledge base ANATOM:

nervous_system[has@(struct) $\rightarrow\rightarrow$ {cns,pns}].

cns[has@(struct) $\rightarrow\rightarrow$ {brain, spinal cord}].

20 brain[has@(struct) $\rightarrow\rightarrow$ {telencephalon, diencephalon,
mesencephalon, rhombencephalon}].

...

cerebellar_cortex[has@(func) $\rightarrow\rightarrow$ {hemisphere, vermis,
flocculus, parafloccular lobes}]

25

eukaryotic_cell :: cell.

brain_cell :: eukaryotic_cell.

neuron :: brain_cell.

glia :: brain_cell, projection_neuron :: neuron.

30

interneuron :: neuron.

purkinje_cell :: projection_neuron.

...

schwan_cell :: glia.

5 The first two groups of facts describe *has-a* relationships which are either structural or functional, the third group specifies the *is-a* hierarchy of brain cells using F-logic's subclass connective " :: ".

Based on such a fact base, rules are used to specify intensional knowledge (which is derivable on-demand at run-time). For example, the
10 transitive closure of *all* has-a relationships (structural, functional, ...) can be expressed by a *single* recursive rule:

$$X[\text{has}@(\text{P}) \rightarrow \rightarrow \{Y\}] \text{ IF } X.\text{has}@(\text{P}).\text{has}@(\text{P}) = Y.$$

thereby illustrating the expressive power of rules using parameterized methods, path expressions, and recursion.

15

Example 7.3 Integration Across Multiple Worlds: Motivating Examples

Consider three research groups: Group A studies neuroanatomy of rodents, group B studies calcium regulatory proteins in vertebrates and group C studies neurotransmission phenomena in mammals.

20

Group A ("Neuroanatomy"). Let us assume that group A has a database of studies, where a study consists of a number of experiments on a specific research problem. Each experiment records a number of experimental parameters in the form of (*name,value*) pairs, and produces a number of
25 images. We focus on a specific image class called *protein labeling images*. For each protein label image the anatomical parameters (i.e., which anatomical region the image represents) and the protein used are recorded. Each image is segmented into a number of segments, based upon the amount of protein staining. The image is also represented as a collection of named anatomical
30 structures visible in the image.

Each anatomical structure is modeled as a collection of segments, such that aggregate features like the distribution of stain within an anatomical structure may be computed. Very often a single biological study involves a number of experiments conducted at *different granularity levels* in the animal.

5 For example, experimenters may try to localize a protein in a tissue, a cell, specific cellular compartments and in intracellular substructures. In this case, the anatomical parameters of an image at any level are *semantically* related to those of an image at the next coarser level although this relation may not be directly visible from the schemas. We will explain how this is modeled
10 in Section 6. In some experiments, specific anatomical structures from a stack of confocal images or a series of electron micrographs are re-constructed into volumetric objects modeled as geometric entities, and specific 3D properties such as the surface to volume ratio are measured. The volumetric information is stored in a separate database DENDREC. The 3D anatomical models are
15 related to the images from which the reconstruction was made. In our example, DENDREC contains the reconstruction of spiny dendrites in the rat neostriatum.

Group B ("Calcium-Binding Proteins"). Next, let us assume group B to have a
20 database of calcium-binding proteins 5 where each protein is identified by its reference number in the PDB 6 and/or the reference number in the SWISS-PROT 7 database. Otherwise it is identified by an internal identifier. A protein has a molecular weight, an amino acid sequence, the number of amino acids and is grouped with a number of other proteins that belong to the same family.
25 Its isoforms, mutants and the species in which the mutants are found are also recorded. Every protein subfamily and mutant form is also given a unique identifier. For every protein the researchers also record the interaction of the protein with elements and ions, the evidence of signal transduction pathways it participates in and the disease processes it contributes to. The database
30 organizes the signal transduction and disease information by the species where the evidence has been found. Also grouped by species, the researchers record

the tissue and cell-level localization they have found in their experiments. However this group does not conduct any experiments at a subcellular level. Although the system does not store the genetic code of the proteins they study, they maintain the reference identifier for the protein form in the GENBANK 8.

5

Group C ("Neurotransmission"). Finally, group C stores information about neurotransmission including neurotransmitter substances, neurotransmitter receptors and voltage-gated conductances in a database NTRANS 9. In this system every neuron is modeled to be composed of a canonical set of non-overlapping compartments. For every compartment of each neuron studied, the experimenters record the input receptors and their description, the intrinsic ionic currents along with their description, and the output transmitters. The description contains a textual account of the function of the receptor or transmitter, the brain region where they are active. Each type of current is characterized by the ions that generate them, their electrical properties, and their firing characteristics. Receptors and transmitters are also organized into families, representable by a tree structure.

10

15

When these systems are integrated, a biologist would like to make queries such as:

20

- *Find the cerebellar distribution of rat proteins with more than 90% amino acid homology with the human NCS-1 protein. Compare the distribution of this protein or its homologs in other rodents.*
- *Are any calcium-binding proteins found only in the thin dendritic spines of the rat neostriatum and not in the stubby spines? Do these proteins always co-localize?*
- *Is there any experiment performed on other mammals on the proteins involved in signal transduction in the visual systems of primates? How similar are these proteins?*

25

A major challenge in today's bioinformatics is to find ways to correlate, combine and unify information from multiple data sources such as described above. But even with many online data sources and information retrieval tools,

30

biologists have little or incomplete technological framework to make this unification possible across disciplines, scales of observation and diversity of viewpoints. As a result, they perform the "integration" task manually, by physically assembling data from multiple sources and putting them together by individual effort. Hence they are seldom allowed the luxury to make "inter-database" queries although the capability to perform such queries is fundamental to the task of the broad-based knowledge unification that they seek. An ostensible source of difficulty arises from the semantic incompatibilities both within and between the data sources.

For example, consider that group A has a number of experiments on the protein distribution of basket cells and their neighboring Purkinje cells without stating that they both belong to the rat cerebellum. We need to model and use this additional piece of knowledge in order to answer the first query. Similarly, the fact that proteins are related because they share amino acid homologies is never recorded because it is "common knowledge" to the domain. However, unless this information is explicitly available from a supplementary source, the query cannot be answered.

In addition to the need for having additional knowledge integration of biological information also have the following issues:

- The information representation at the mediator should be flexible enough to accommodate a wide degree of heterogeneity at the data sources, and at the same time, represent the class-structure evident from the taxonomic character of the data. To accommodate this, we use an object-oriented formalism, but unlike the collection-based model used in the prior art, we use F-logic that is well-equipped to represent object-orientation, flexible enough to represent semistructured data, and has the machinery to perform inferences and recursive computation such as path expressions and transitive closure.
- The computation of numeric aggregates and numeric features describing the content of 2D and 3D images and reconstructed volumes is an essential component of the data to be integrated. Equally important is

the need to represent complex semantic rules to model the associations between numeric features computed from multiple image and volume instances.

- Queries involving graph operations such as graph intersection and computation of the spanning tree are important in discovering the associations between data coming from different sources that are initially unconnected. Meta-level reasoning with schema and attributes are an important component in creating these associations.

10 Example 7.4 The Integration Framework

Most current approaches to integration of information from heterogeneous sources are based on the prior art mediator architecture as discussed herein above in the Background of the Invention section. The problem of heterogeneous data models of sources is solved by translating the data into a common language using *wrappers*. The *semistructured data model* (essentially *labeled directed graphs*) in general and XML in particular have been shown to be suitable target data models. Once the data can be accessed in a uniform way, a *mediator* is used to integrate between the different local views and schema elements, based on the specification of an integrated view.

The definition of such an integrated view can often be a highly complex task and requires dealing with all of the well-known integration problems from information integration in databases like structural, semantic, and descriptive conflicts (e.g., flat vs. nested relational vs. object-oriented modeling, homonyms, synonyms). Thus, for complex integration tasks, a powerful declarative specification language is required, e.g., for querying and restructuring local schemas, mapping data between models and schemas, integrity checking, and knowledge inference.

Example 7.4.1 The KIND Architecture

For our neuroscience application domain, an invention embodiment called KIND, with the schematic of FIG. 9 illustrating an architecture of the

embodiment. The main data sources are scientific STUDIES of various types which themselves can refer to further heterogeneous data sources like PROLAB (image databases of protein labelings) and DENDREC (volumetric reconstructions of dendrites). Other KIND sources are CAPROT (calcium-binding protein databases) and NTRANS (neuro-transmission database). Apart from these sources of observational data, there are also sources with general domain data and knowledge like ANATOM (anatomical knowledge base) and TAXON (animal taxonomy database).

Unlike prior art mediator approaches that solely use the semistructured model throughout the integration, the present embodiment additionally incorporates a rich object-oriented knowledge representation formalism, i.e., F-logic into the architecture.

This enables a better modularization and more adequate modeling of complex application domains like biology and neuroscience. In our biological integration domain, for example, links between otherwise unrelated data are established using expert knowledge (like anatomical, taxonomic, or partonomic relationships) that is represented using F-logic rules. In some cases, ontologies exist for modeling specific aspects or parts of a domain, thereby providing a unique semantics for that part. While an ontology captures the semantics of some domains, the problem remains to mediate across different ontologies for providing the user with an integrated view. Again a powerful integration language like F-logic is needed for mediating between the ontologies.

Example 7.4.2 KIND Modules

The *source modules* of the KIND architecture, i.e., data and knowledge sources, have an associated XML DTD (*Document Type Definition*) describing the structure of the exported data after wrapping. Here, we speak of a *data source* (or data module), when the modeled information has mainly observational character like data collected during an experiment, and of a *knowledge source* when we model information about the application domain ("general" or "expert knowledge"), usually in rule form. Note that for some

sources wrapping may be done once and off-line. However, often this translation has to be done online (or *on-demand*), i.e., the wrapper has to translate incoming XML queries to native queries against the actual source data. In general, the *query capabilities* of the underlying source are limited in which case the wrapper can support only specific XML queries.

Syntactic integration with some minimal consistency at the source level is achieved by enforcing that a source module M exports *valid* XML, i.e., which conforms to the associated $DTD(M)$. Clearly, additional integrity constraints dealing with both structural aspects and application domain constraints should be modeled in order to guarantee consistency at a higher, conceptual level.

Example 7.4.3 Exported Object and Class Structure

At the level of a source module M , we incorporate F-logic by providing a class signature $\Sigma(M)$ of exported classes and their objects' structure, which constitutes a semantically much richer conceptual-level specification of the source than just $DTD(M)$ and thereby facilitates the integration of M at a conceptual level at the mediator. In particular, $\Sigma(M)$ specifies:

- the source's *class hierarchy*,
- whether attributes (i.e., F-logic methods) are *single-valued* or *multi-valued*,
- whether and how they are *parameterized*, and
- whether they are *inheritable* or not.

Formally, the object-oriented class structure of M is given by a mapping Φ_M : $DTD(M) \rightarrow \Sigma(M)$. Technically, Φ_M is straightforward to implement as it amounts to a simple syntactic transformation from XML elements to F-logic expressions (e.g., using an XML parser whose output is "pretty-printed" to F-logic, or using the XML stylesheet/transformation language XSL(T)). The difficulty consists in choosing the most appropriate "semantically adequate" representation in F-logic of the underlying, XML-encoded, object model.

Consider, for example, a generic mapping Φ_{GEN} which maps *arbitrary* XML documents (i.e., irrespective of an object model of the encoded information) to F-logic representations: Since an XML document is a semistructured database (more precisely, a labeled ordered tree) it can be
 5 represented in F-logic, for example, over the signature XML node:

```
[element_type  $\Rightarrow$  string;
  attribute@(string)  $\Rightarrow$  string;
  child@(integer)  $\Rightarrow$  xml_node ].
```

While Φ_{GEN} faithfully represents any given XML document, the
 10 application domain structure is *not* visible at the schema level and has to be extracted from the data. Thus, whenever possible, it is preferable to model a source by first specifying its application domain structure in F-logic, i.e., designing $\Sigma(M)$. Then a syntactic representation of $\Sigma(M)$ using an XML DTD is straightforward, and we can trivially go back from that DTD to $\Sigma(M)$. Hence
 15 we get Φ_M essentially "for free".

In case a source module does not have an F-logic signature $\Sigma(M)$, for example, because M is a new source module being added to the system and $\text{DTD}(M)$ is unknown, or Φ_M has not yet been established, then Φ_{GEN} can
 still be useful as a first means to bring the new data into the system. Indeed, F-
 20 logic is also suitable as a language for managing semistructured data, i.e., extracting data using generalized path expressions, discovering schema etc.

Example 7.4.4 Exported Integrity Constraints

In addition to $\Sigma(M)$, a set of application specific *integrity constraints*
 25 $\text{IC}(M)$ can be provided. These are F-logic rules that create "alerter objects", i.e., instances of class alert whenever an inconsistency (at the class or object level) is derived. An alerter object indicates the type of inconsistency encountered and some hints on which objects and classes were involved in the inconsistency which greatly simplifies debugging the data. In particular, this

allows to differentiate between *local inconsistencies* (i.e., within a module M) and *global inconsistencies*.

Example 7.4.5 Derived Knowledge

- 5 Finally, some modules M also export *intensional knowledge* in the form of a set of F-logic rules IDB(M).

Example 7.4.6 ANATOM Fragment Cont'd

- 10 The anatomical knowledge base ANATOM referenced above includes the following rules:

purkinje_cell[located_in →→ {purkinje_cell_layer}].
 basket_cell[located_in →→ {cerebellar cortex}].
 X[located_in →→ C] IF
 15 X: nucleus[located_in →→ {N: neuron}],
 N[compartments →→ {C: cell_body}].
 X[located_in →→ {X}] IF X: neuro_anatomic_entity.
 Y[located_in →→ {X}] IF X[has@(P) →→ {Y}].
 X[located_in →→ {Y}] IF X..located_in..located_in = Y.

20

defining the located_in relation from base facts, specific anatomical knowledge rules, and generic rules for defining reflexive and transitive closure.

Example 7.4.7 The KIND Mediator

- 25 As explained above, the structure and semantic information of a source module M is specified using an XML DTD(M) (mainly for inputting the wrapped raw data), a class signature $\Sigma(M)$, the correspondence mapping Φ_M between them, integrity constraints IC(M) and, in the case of derived knowledge, IDB(M). The KIND mediator module itself exports an integrated
 30 F-logic view INSM (*Integrated NeuroScienceModel*) to the user, which is

defined based on the imported source modules (STUDIES, TAXON, ...), the facts and rules from the imported knowledge bases (ANATOM, ICORR), and the actual view-defining integration rules. The mediator imports signatures of a module M using declarations like:

```

5      import study[
          id⇒string; project⇒string;
          experiments⇒⇒experiment; ...]
      from 'STUDIES'.

```

In this way, a subgoal of the form $S: \text{study}[\text{id} \rightarrow I; \text{project} \rightarrow P; \dots]$ induces a query against the STUDIES source. Note that oids of objects from different source modules M1 and M2 are guaranteed to be distinct. The declarative way to achieve this is by qualifying each oid with the URI of the module from which it was imported (in the implementation we can just use disjoint sets of integers as oids). The only oids which can be shared across modules are those of string objects and constants occurring in the import declaration. For example, consider two sources SIMPLE and DETAILED of animal data. We can simultaneously import from both modules as follows:

```

      :- import animal[name⇒string] from 'SIMPLE'.
      :- import animal[common name⇒string; species⇒string;
20  genus⇒string; ...] from 'DETAILED'.

```

Constants appearing in the import declarations (animal, species, : : :) and string-valued objects like "Rodentia" are distinguished and thus shared external object names. Therefore a subgoal of the form $X: \text{animal}[\text{Attr} \rightarrow \text{Val}]$ will yield both, instance from SIMPLE and from DETAILED, together with their attribute/value pairs. As part of the integration process, we may have to distinguish between instances of animal from SIMPLE and those from DETAILED (note that the way internal oids are differentiated may not be visible to the rule programmer).

This is accomplished by qualifying names with the module they were important from: e.g., $X: ('DETAILED'.\text{animal})(M \rightarrow R)$ will only range over

objects from DETAILED. Logically, this corresponds to defining for each module M , the methods $M.N$ for all distinguished names imported from M :

$M[N \rightarrow M.N]$ IF N : distinguished name[imported_from $\rightarrow M$].

When importing data from M , a distinguished (exported) class name C is
5 prefixed with M and all instances of C in M are made instances of $M.C$.

Example 7.4.8 F-Logic Query Evaluation

The invention embodiment of Example 7's KIND program and method uses a central mediator component with FLORA, an F-logic to XSB-Prolog
10 compiler, as the evaluation engine. More information regarding FLORA may be had at "FLORA: The Secret of Object Oriented Logic Programming", Ludascher, Yang and Kifer, Technical Report, State University of New York, Stony Brook, June 1999, incorporated herein by reference. Due to its built-in
top-down strategy, FLORA derives facts in a *demand-driven* way somewhat
15 similar to the VXD architecture of MIX_m. At the current implementation stage, sources have no independent query evaluation mechanism but simply export all data and rules to the central mediator. However, the design of the architecture allows for source modules to have their own evaluation engine in which case
source data is imported only as needed for answering queries.

20 In contrast, the FLORID system is an implementation of F-logic which employs a bottom-up and thus a model *materialization* strategy. More information regarding FLORID may be had at "Managing Semi-structured Data with FLORID: A Deductive Object-Oriented Perspective", Ludascher, Himmeroder, Lausen, May, and Schlepphorst, Information Systems, 23(8):589-
25 613, 1998; herein incorporated by reference. It has been shown that FLORID is well-suited for management of semistructured data and as a unified framework for wrapping and mediating Web data. Therefore, it may be desirable to incorporate the FLORID engine into the KIND architecture for modules and
views where materialization is advantageous.

30

Example 7.5: Elements of the Integrated Neuroscience Model

In this section concepts described in the previous section are illustrated by examples.

Example 7.5.1 XML-DTD and F-logic Representations

- 5 Each source module *M* has an associated XML DTD. The XML data may result from wrapping of the raw data, or the source may natively support XML. The mediator can either import the XML DTD as is (using the generic mapping *_M* to F-logic), in which case any application specific structure not visible from the DTD has to be "re-covered" at the mediator, or the mediator
10 can import the semantically richer F-logic signature.

The following XML DTD is used by the STUDIES database:

```

<!ELEMENT Studies (Study)*>
<!ELEMENT Study (study_id, project_name,
    project_description, animal, experiments, experimenters)>
15 <!ELEMENT animal (subject_id, scientific_name, strain, age)>
    <!ELEMENT experiments (experiment)*>
    <!ELEMENT experiment (description, instrument, parameters)>
    <!ELEMENT instrument (type, name)>
    <!ELEMENT parameters (parameter)*>
    <!ELEMENT parameter (name, value)>
20 <!ELEMENT experimenters (experimenter)*>
    <!ELEMENT experimenter (name, affiliation)>

```

One of many ways to model this in F-logic is as follows:

```

25 studyDB[studies⇒⇒study].
   study[id⇒string; project name⇒string; description⇒string;
   animal⇒animal; experiments⇒⇒experiment;
   experimenters⇒⇒string].
   animal[subject id⇒string; scientific name⇒string; strain⇒string;
30 age⇒string].
   experiment[description⇒string; instrument⇒instrument;

```

parameters $\Rightarrow\Rightarrow$ exp_parameter].

instrument[type \Rightarrow string; name \Rightarrow string].

exp parameter[name \Rightarrow string; value \Rightarrow string].

In general, the F-logic signature $\Sigma(M)$ can carry much more semantic
 5 information from the application domain (due to class hierarchies,
 parameterized methods, single- vs. multi-valued, etc.), in particular when
 accompanied by integrity constraints $IC(M)$ and derived knowledge $IDB(M)$.

Example 7.5.2 Creation of Mediated Classes

10 At the mediator level, the F-logic schema of the source is modified to
 relate it to the knowledge sources. For example, the type of scientific name,
 which consists of genus and species names and may optionally include a
 subspecies name, is modified from a string to a taxon reference, where a taxon
 is an element from the taxonomic database. The modification is made by first
 15 creating a new class called animal at the mediator as the union of the classes
 called animal at the sources. Thus, at the mediator

animal[M \Rightarrow R] IF S: source, S.animal[M \Rightarrow R].

Then a new method is added to this union class to link it to the taxonomic
 database:

20 animal[taxon \Rightarrow 'TAXON'.taxon].

Finally, the association between the scientific name in the PROLAB database
 and the taxonomic database is created:

X[taxon \rightarrow T] IF : 'PROLAB'.animal[scientific_name \rightarrow N],
 words(N,[W1,W2j]), T : 'TAXON'.taxon[genus \rightarrow W1;
 25 species \rightarrow W2].

The built-in predicate words, when given a string as first argument,
 returns the list of words of that string. As will be used later, this predicate can
 also create a string of words which are separated by a whitespace. Such
 somewhat "procedural" predicates like word can easily be defined at the

mediator, since the whole XSB-Prolog machinery is accessible from the FLORA F-logic engine.

Example 7.5.3 Geometric Modeling

- 5 In modeling reconstructed volumes of dendritic spines [17], we first create a number of solid-geometric primitives called `shape3D` like cylinders, spheres and hyperboloids, used by solid modeling software.

```
shape3D[volume⇒scalar, area⇒scalar]
```

```
cylinder :: shape3D[radius⇒scalar, length⇒scalar].
```

- 10 The model of dendritic spines is composed of these primitives:

```
spine :: shape3D[view files⇒⇒url].
```

```
mushroom spine :: spine[head⇒head; taper⇒taper, neck⇒neck].
```

```
head[shape⇒sphere]. taper[shape⇒hyperboloid].
```

```
neck[shape⇒cylinder].
```

- 15 The `view files` attribute yields a list of urls that represent the images of different projections of the 3D volume. A parametric attribute of F-logic is used to model spines protruding out from the shaft of a dendrite at a coordinate (x,y,z):

```
shaft[connected spine@(x,y,z)⇒spine; num spines⇒integer].
```

- 20 **Example 7.5.4 Rules for Classification and Integrity**

The dendritic spines are classified into thin, stubby and mushroom classes using F-logic rules:

```
S: mushroom spine IF S: spine[head→_ ; neck→_ ; taper→_ ].
```

```
S: stubby spine IF S : spine[head→_ ; undef→→{neck, taper}].
```

- 25 S: thin spine IF S: spine[neck→_ ; undef→→{head, taper}].

The method `undef` applied to an object `O` yields those methods `M` that are declared for class `C` but which are not defined for `O`. This is a simple example for reasoning about schema and is specified in F-logic as follows:

```
O[undef→→{M} IF O: C[M→→_ ], not O[M→→_ ].
```

The dendritic reconstruction data source also has integrity constraint rules, defined as a special class called alert. For example, the constraint that a dendritic spine cannot have only a taper (but no head or neck) is modeled as:

ic1(S) : alert[type \rightarrow "singleton taper"; object \rightarrow S] IF S:
5 spine[taper \rightarrow _ ; undef \rightarrow \rightarrow {head,neck}].

Example 7.5.5 Complex Relationships

As mentioned earlier, experimental biological information often have complex semantic relationships. For example, two experiments in a single
10 study in source A may be related in the following way. In the first experiment, the experimenters perform a protein labeling on the entire brain and record the result as a segmentable image. In the second experiment, they would like to investigate the protein labeling pattern of the heavily stained portions of the brain region called *cerebellum*. So they extract that part of the cerebellum
15 (from an identical specimen) which showed heavy staining in the previous experiment, and produce a finer resolution image to identify the actual cells that took the heavy stain. Although these two images are related, the relationship cannot be modeled just by linking the second image to a segment of the first. In reality, the second image is related to *any segment* in the first
20 specimen that satisfies the condition of being "heavily stained" and in the cerebellum. We model this by using *named predicates*.

Consider a fragment on the schema of the class image:

image[anatomical structures \Rightarrow anatomical structure].
anatomical structure[name \Rightarrow string; segments \Rightarrow segment].
25 segment[description \Rightarrow string; features \Rightarrow feature].

We consider the simple case where a feature is a single floating point number. Let us assume that the class protein_label_image :: image has only one feature called protein amount. In order to express that a segment is *heavily stained* we can specify a user-defined predicate has prop as follows:

30 has prop(I, heavily stained, S) IF S.features.protein amount > 100 .

meaning that segment S of image I has the property "heavily stained" if the staining intensity is greater than 100. We assume a relation derived with (P, I_1, I_2) that is instantiated every time a researcher creates a finer resolution image I_2 based upon some property P on a coarser resolution image I_1 . Then the rule:

```

5      highlight_parent_segments(I2,S) IF
        derived with (P, I1,I2), I1 : protein label image, I1[anatomical
        structures..segments→→{S}], has prop(I2,P,S).

```

can be used to encode the relationship between the two images. Used this way, the rule will produce all possible segments in I_1 that could have produced the image I_2 . We could also use a rule with the same body to return all derived
10 images from a given image.

Example 7.5.6 Meta-Reasoning with Schema

15 The mediator performs meta-reasoning of the schema of TAXON to create a class hierarchy of animals. Consider the schema of TAXON:

```
taxon[subspecies⇒string; species⇒string; genus⇒string;
      family⇒string; order⇒string; infraclass⇒string, ...
      ... phylum⇒string; kingdom⇒string; superkingdom⇒string].
```

At the mediator, a hierarchy is defined for the taxonomic ranks:

20 subspecies :: species :: genus :: ... :: kingdom :: superkingdom.

Now the *data* in the TAXON database is used to infer the taxonomic *class hierarchy*:

T : TR, TR:: TR1 IF

25 **T : 'TAXON'.taxon[Taxon Rank→TR; Taxon Rank1→TR1],**
 Taxon Rank :: Taxon Rank1.

The rule states that given two taxon ranks, e.g., order and kingdom with data values *rodentia* and *metazoa* respectively, and given that kingdom is a subclass of class, then *rodentia* is a subclass of *metazoa*. In other words, from the *data* of TAXON we infer new *schema information*, i.e., that all rodents belong to the metazoe kingdom. As we will show in the next section, this rule

will be used in a query to determine the appropriate taxonomic ranks for computing joins and closures.

Example 7.5.7 Rule Export from Knowledge Bases

5 Knowledge bases export rules to the mediator. The anatomical knowledge base, for example, contains both an "is a" and a "has a" hierarchy. Thus a Purkinje cell "is a" neuron and cerebellum "has a" Purkinje cell layer. We also use the predicate "located in" as an inverse of the "has a" relation. Thus the fact that Purkinje cell is "located in" Purkinje cell layer implies it is
10 also "located in" the cerebellum. This rule is used in the mediator to create a transitive closure over the locations of neuro anatomic entity(ies) during a query.

Example 7.6 Semantic Integration from the Mediator's and User's Perspective

15 To illustrate how an integrated query is evaluated in the Integrated Neuroscience Model embodiment of the invention, we trace through the phases of evaluating the first example query:

(1) *Find the cerebellar distribution of all rat proteins with more than 90% amino acid homology with the humanNCS-1 protein.*

20 The broad steps for evaluating this query in the INSM module are: (i) retrieve facts about *shared homologies* where homology>90% (uses CAPROT), (ii) determine the *protein distribution* using data from PROLAB and ANATOM, and (iii) compute the *aggregate*, grouped by anatomical structure.

More precisely, let us assume that the mediator defines and exports the
25 following two views called *homologous_proteins* and *aggregated_protein_distribution*. The first can be treated as a relation:

Homologous_proteins(Protein1,Animal1, Protein2, Animal2,
Name_type, Value)

Here the two protein-animal pairs refer to the variety of the specified protein as
30 found in the given animal. This relation depicts that given two such pairs, the database stores how similar they are in terms of their amino acid sequence as a

percentage value. The attribute name type specifies whether the common name or the scientific name of the animals have been specified in the query 20. Similarly, the second view can be treated as the relation

5 aggregated_protein_distribution(Protein, Organism name, Name_type,
Brain_region, Feature_name, Anatom_struct, Result)

The relation records the distribution of a feature (such as protein amount) of proteins occurring in the brain region of an organism, grouped by the anatomical structures in that brain region. The user's query is expressed in terms of these

10 views in the following manner.

query1(Anatom_struct, Result) IF homologous proteins("NCS-1",
"human", Rat Protein, "house rat", common, Value),
Value > 90,
15 aggregated_protein_distribution(Rat Protein, "house rat", common,
"cerebellum", "protein amount", Anatom_struct, Result).

In the mediator the first view homologous proteins is constructed by importing from module CAPROT the class amino acid homology. The view definition of homologous proteins based on amino acid homology is:

20 homologous_proteins(Protein1,Animal1, Protein2, Animal2, Name type,
Value) IF
:_amino acid homology[shared@(
_: protein in animal[
name!Protein1;
found in→_: animal[
25 name@(Name type)→Animal1]],
_: protein in animal[name→Protein2;
found in→_: animal[
name@(Name type)→Animal2]])
→Value].

Note that since this entire view is in the scope of one source module we do not need to qualify names by the module name. The definition of the second view illustrates the use of aggregation (here: summation of Values, grouped by Anatom struct):

```

5      aggregated protein distribution(
        Protein, Organism name, Name type, Brain region, Feature name,
        Anatom struct, Result) IF
        Result = sum{Value [Anatom_struct] ;
10      protein_distribution(Protein, Organism_name
        Name_type, Brain_region, Feature_name,
        Anatom_struct, Value)}.

```

Here the view protein distribution is defined by importing the protein label image class of module PROLAB and the class neuro anatomic entity class of module ANATOM. Finally, a *semantic join* based on the ANATOM

15 knowledge base is illustrated by the following rule:

```

        protein_distribution(Protein, Organism_name, Name_type,
        Brain_region, Feature_name, Anatom_struct, Value) IF
        I:'PROLAB'.protein label image[proteins→→Protein; organism@
        Name_type)→Organism_name; anatomical_structures→→
20      {A:'PROLAB'.anatomical_structure[name→Anatom_name]}],
        NAE: 'ANATOM'.neuro_anatomic_entity[name→Anatom_name;
        located_in→→{Brain_region}],
        A..segments..features[name→Feature_name; value→Value].

```

In this view-definition rule the last two arguments of protein distribution are
25 used as output variables while the rest are used as input variables. The anatomical structure from the PROLAB and ANATOM modules are explicitly joined using the variable Anatom name. As explained before, the recursive definition of located in in the ANATOM module, causes the rule to transitively traverse every substructure of the cerebellum down to the cellular level in order
30 to find the "leaf level" anatomic structures where the protein is localized. This

constitutes the semantic join between Brain region and the anatomical structure A whose features are being extracted.

(2) Compare this with the distribution of this protein or its homologs in other rodents.

The primary difference between this query and the previous one is that it is executed over the set of *all rodents except rat*, and that it uses information from the module TAXON. With this modification, the second query is stated as:

```
10      query2(Anatom_struct, Result ) IF homologous_proteins("NCS-1",  
      "Homo sapiens", Rodent_Protein, Rodent_name, scientific, Value), :  
      'TAXON'.taxon[order→"Rodentia"; genus→G; species→S ],  
      words(Rodent_name, [G,S]), Scientific name = \= "Rattus rattus",  
      Value > 90,  
15      aggregated_protein_distribution(Rat Protein, "house rat",  
      common, "cerebellum", "protein_amount",  
      Anatom_struct, Result).
```

In this query we use the scientific rather than the common names of organisms, and we explicitly use the information that "Rodentia" is a value of the *order* attribute of the class taxon. This directly collects all known rodents in a set over which the rest of the query is evaluated. A less straightforward (but perhaps easier for the less knowledgeable user) way of evaluating the query could be to walk the transitive relationship of taxonomic classes to discover that we need all species under the order "Rodentia". Also note that the species "Rattus rattus" has been explicitly eliminated from the set to compute the rest of the aggregated protein distributions:

Those knowledgeable in the art will appreciate that the invention as claimed in the claims appended hereto is not limited to the embodiments described and discussed herein, and that many equivalent elements and steps may be substituted for those used herein. For example, many equivalents to the

particular languages used herein such as XML, F-logic, RDF, UML, EER, and the like are of course possible. Also, the invention is not limited to practice on databases, but is instead able to be practiced with any of a variety of data sources that can be queried, with examples including, but not limited to, databases, tables, spreadsheets, forms, web pages with forms, flat files, software with API's (application program interfaces), and the like. The invention is also not limited to practice with any particular number of data sources. Indeed, although examples herein have referred to practice with two or three data sources, those knowledgeable in the art will appreciate that the invention may be practiced with multiplicities of different data sources.

The present invention thereby provides a novel and elegant solution to problems that were heretofore unresolved. An embodiment of the invention generally comprises integrating data sources at a "knowledge" or "conceptual" level that sits above the "structural level". The higher level integration allows for deeper and richer relations between data sources to be manipulated for resolving queries. Also, an invention embodiment provides a highly extendable architecture allowing for easy addition of data sources to the integration regardless of the source's formalism through use of a generic conceptual model and "plug ins" for converting between the source's formalism and the generic model formalism. Still further, an invention embodiment allows for responsive data to be presented in context by displaying not only the data but also the logic relations that were made to relate classes in obtaining the data.

The advantages of the disclosed invention are thus attained in an economical, practical, and facile manner. While preferred embodiments and example configurations have been shown and described, it is to be understood that various further modifications and additional configurations will be apparent to those skilled in the art. It is intended that the specific embodiments and configurations herein disclosed are illustrative of the preferred and best modes for practicing the invention, and should not be interpreted as limitations on the scope of the invention as defined by the appended claims.

What is claimed is:

1. A method for integrating a plurality of data sources comprising the steps of:
 - 5 obtaining semantic information from each of the plurality of data sources;
 - creating a conceptual model for each of the plurality of data sources using said semantic information;
 - accessing a secondary knowledge source having information that
 - 10 relates the different data sources;
 - creating an integrated semantic model of the plurality of data sources using said conceptual models and said secondary knowledge source.
2. A method as in claim 1 wherein each of said conceptual models
- 15 is in a uniform language.
3. A method as in claim 1 wherein each of the plurality of data sources has a semantic formalism, and wherein the method further comprises providing a plug-in for converting at least one of the plurality of data source's
- 20 semantic formalism to said uniform language of said generic conceptual model.
4. A method as in claim 1 wherein said semantic information obtained from at least one of the plurality of data sources is expressed in one of either a unified modeling language, resource description framework, or an
- 25 extended entity relationship language.
5. A method as in claim 1 wherein the step of creating a conceptual model for each of the plurality of databases further comprises converting each of said conceptual models into a generic conceptual model, each of said generic
- 30 conceptual models being in a uniform language.

6. A method as in claim 5 wherein said uniform language represents logic statements.

7. A method as in claim 5 wherein the step of creating said
5 integrated semantic model comprises using said generic conceptual models.

8. A method as in claim 1 wherein each of said conceptual models created from each of the plurality of data sources comprise logic statements, and wherein said integrated semantic model comprises logic statements.
10

9. A method as in claim 1 wherein said secondary knowledge source comprises logic statements in a frame logic language.

10. A method as in claim 1 wherein said integrated semantic model
15 comprises logic statements, and wherein at least one of said logic statements comprises an instance declaration, at least one of said logic statements comprises a subclass declaration, and at least one of said logic statements comprises a method declaration.

20 11. A method as in claim 1 wherein said secondary knowledge source comprises a plurality of secondary knowledge sources.

12. A method as in claim 1 wherein each of the conceptual models from each of the plurality of data sources have at least a class, said semantic
25 information comprises at least a class schema, and wherein said integrated semantic model comprises logic statements that define one of at least an instance, subclass, or method relation between at least a first class from one of the plurality of data sources and a second class from a second of the plurality of data sources.

13. A method as in claim 1 wherein each of said conceptual models has at least a class, said semantic information comprises a class schema, and wherein said integrated semantic model comprises logic statements that define an instance, subclass, and a method relation between one or more classes from
5 a first of the plurality of data sources and one or more classes from others of the plurality of data sources.

14. A method as in claim 1 wherein each of said conceptual models has at least a class, and wherein said integrated semantic model comprises at
10 least a logic statement that defines one of an "is a", a "has a", or a "part of" relation between at least a first class from one of the plurality of data sources and one or more classes from others of the plurality of data sources.

15. A method as in claim 1 wherein said secondary knowledge
15 source comprises expert knowledge.

16. A method as in claim 1 wherein each of the plurality of data sources define at least a class, and wherein said secondary knowledge source has information comprising rules that relate at least some of the classes from
20 different of the plurality of data sources to one another.

17. A method as in claim 16 wherein said rules relate at least a first class from one of the plurality of data sources to one or more second classes from others of the plurality of data sources in a logical instance, subclass, or
25 method relation.

18. A method as in claim 1 wherein said secondary knowledge source information is not available from any of the plurality of data sources.

19. A method as in claim 1 wherein said method further comprises converting data from each of the plurality of data sources into a uniform language.

5 20. A method as in claim 19 wherein said uniform language comprises extensible markup language.

21. A method as in claim 1 wherein the method further comprises processing a user-input query against said integrated semantic model to identify
10 data from one or more of the plurality of data sources responsive to said query.

22. A method as in claim 21 wherein each of the plurality of data sources comprises at least a data class, and wherein in responding to said query the method further comprises the steps of:
15 logically relating one or more classes from different of the plurality of databases to one another, and
presenting said data responsive to the query in context by displaying said logic relations made between said related classes.

20 23. A method for integrating a plurality of individual data sources to process a query comprising the steps of:
converting data from each of the plurality of data sources to a uniform language;
obtaining semantic information from each of the plurality of data
25 sources, said semantic information comprising at least a class schema;
creating a generic conceptual model in a uniform language for each of the plurality of data sources using said semantic information, said uniform language comprising logic statements;
accessing at least one secondary knowledge source comprising
30 rules relating data classes from different of the data sources to one another;

creating an integrated semantic model using said conceptual models and said secondary knowledge source, said integrated semantic model comprising logic statements expressing at least one of an instance, subclass, or method relation between a first class from one of the plurality of data sources and one or more second classes from another of the plurality of databases;

converting the query to a language consistent with said integrated semantic model; and

using said integrated semantic model to identify data from one or more of the plurality of data sources responsive to the query.

10

24. A method as in claim 23 wherein each of the plurality of data sources has a semantic formalism, and wherein said data responsive to the query comprises data retrieved from at least two of the plurality of databases, and wherein the method further comprises the steps of:

15 providing a plug-in for converting between at least one of the plurality of data source formalisms and said generic conceptual model uniform language; and

displaying said data responsive to said query in context by displaying logical relations between classes from different of the plurality of data sources made to determine said responsive data.

20

25 25. A computer program for integrating a plurality of different data sources, the computer program comprising computer executable instructions stored in a computer readable medium that when executed cause the computer to:

obtain semantic information from each of the plurality of data sources;

create a conceptual model for each of the plurality of data sources using said semantic information;

30 access a secondary knowledge source having information relating the data sources to one another; and

create an integrated semantic model using said conceptual models and said secondary knowledge source.

26. A computer program as in claim 25 wherein each of said
5 conceptual models comprises a generic conceptual model in a uniform language.

27. A computer program as in claim 26 wherein each of said generic
conceptual models comprise logic statements, and wherein said integrated
10 semantic model comprises logic statements.

28. A computer program product as in claim 26 wherein the
computer program instructions when executed further cause the computer to
provide a plug-in for converting between one of the plurality of data source's
15 semantic formalism and said uniform language of said generic conceptual
model.

29. A computer program as in claim 25 wherein said semantic
information obtained from at least one of the plurality of data sources is
20 expressed in one of either a unified modeling language, resource description
framework, or an extended entity relationship language.

30. A computer program as in claim 25 wherein said semantic
information comprises class schema and relationship schema.
25

31. A computer program as in claim 25 wherein each of the plurality
of databases has at least a class, and wherein said integrated semantic model
comprises logic statements relating at least some of said classes from different
of the data sources to one another.
30

32. A computer program as in claim 31 wherein said logic statements comprise Frame logic.

33. A computer program as in claim 25 wherein said integrated semantic model comprises logic statements relating classes of data from different of the plurality of databases to one another, and wherein at least one of said logic statements comprises an instance declaration, at least one of said logic statements comprises a subclass declaration, and at least one of said logic statements comprises a method declaration.

10

34. A computer program as in claim 25 wherein each of the plurality of data sources has at least a class, and wherein said integrated semantic model comprises logic statements that express an instance, a subclass, and a method relation between at least one class from a first of the plurality of databases and one or more classes from others of the plurality of databases.

15

35. A computer program as in claim 25 wherein each of said conceptual models comprises a class hierarchy having at least a class and a subclass.

20

36. A computer program as in claim 25 wherein said secondary knowledge source comprises expert knowledge.

37. A computer program as in claim 25 wherein said secondary knowledge source comprises a digital data file.

25

38. A computer program as in claim 25 wherein said secondary knowledge source comprises a plurality of secondary knowledge sources.

30

39. A computer program as in claim 25 wherein each of the plurality of data sources define at least a class, and wherein said secondary knowledge

source has information relating at least some of the classes from different of the plurality of data sources to one another.

40. A computer program as in claim 25 wherein said secondary
5 knowledge source information relates at least a first class from one of the plurality of data sources in one of an instance, sub-class, or method relation to a second class from another of the plurality of databases.

41. A method as in claim 25 wherein said secondary knowledge
10 source information is not available from any of the plurality of databases.

42. A computer program product as in claim 25 wherein the program
instructions when executed cause the computer to convert data from each of the
15 plurality of data sources into a uniform language.

43. A computer program product as in claim 42 wherein said uniform
data language comprises a semistructured language.

44. A computer program product as in claim 43 wherein said uniform
20 language comprises extensible markup language.

45. A computer program product as in claim 25 wherein the
computer program instructions when executed cause the computer to process a
user-input query and to use said integrated semantic model to identify data
25 from one or more of the plurality of data sources responsive to said query.

46. A computer program product as in claim 45 wherein the
computer program instructions further cause the computer to display said data
responsive to said query in context by displaying logic relations made between
30 classes from different of the plurality of data sources to identify said responsive
data.

47. A computer program product for querying a plurality of individual heterogeneous data sources, the computer program product comprising computer readable instructions embedded in a computer readable medium that when executed cause the computer to:

obtain data from each of the plurality of data sources in a uniform, semi-structural language;

obtain semantic information from each of the plurality of data sources, said semantic information comprising at least a class schema;

10 convey said semantic information to a mediator;

create a generic conceptual model in a uniform language with said mediator for each of the plurality of data sources using said semantic information;

15 create a domain map with said mediator using at least one secondary knowledge source, said domain map comprising logic statements that at least relate a class for a first of the plurality of data sources to a class from another of the plurality of data sources in one of a instance, class, or method relation; and

20 process a query with said mediator by converting said query to a language consistent with said domain map, using said domain map to identify classes from one or more of the plurality of data sources responsive to said query, retrieve said responsive data, and display said responsive data.

48. A computer program product as in claim 47 wherein each of the plurality of data sources has a semantic information language, and wherein the product further causes the computer to:

provide a plug-in for converting between at least one of the data source semantic information languages and said generic conceptual model uniform language; and

30 display said responsive data in context by displaying logic relations made between said classes to identify said responsive data.

1/7

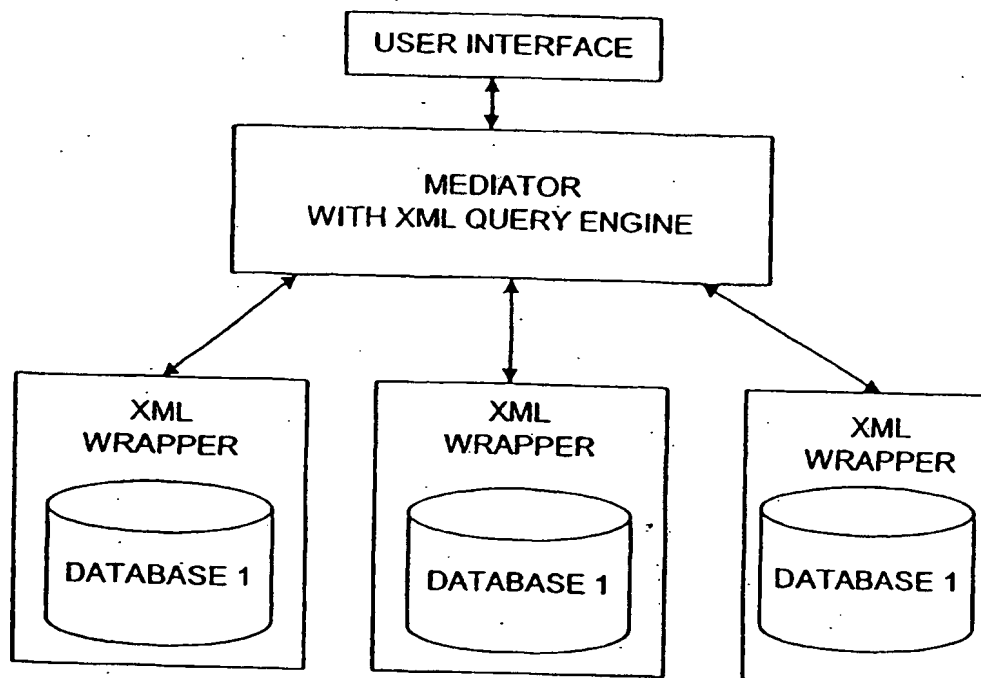
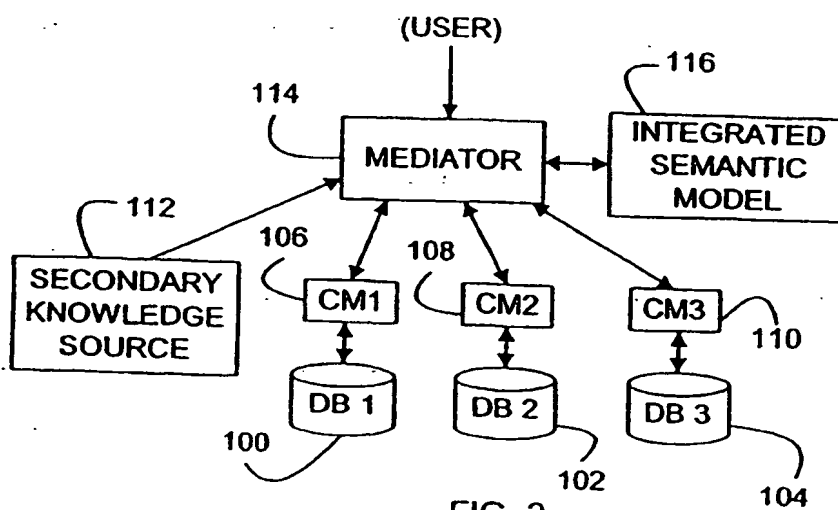
FIG. 1
PRIOR ART

FIG. 2

2/7

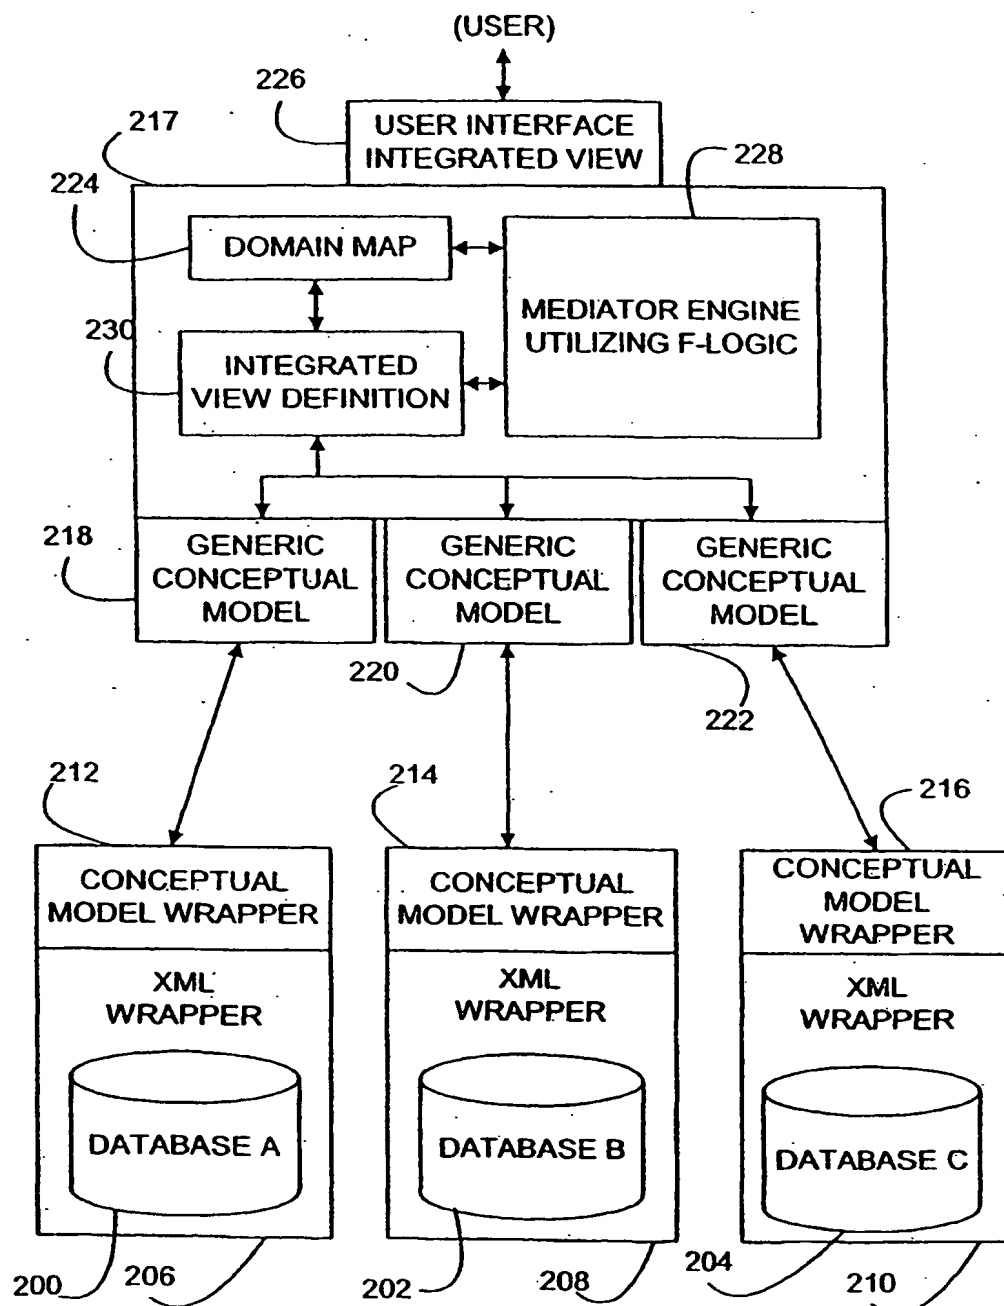


FIG. 3

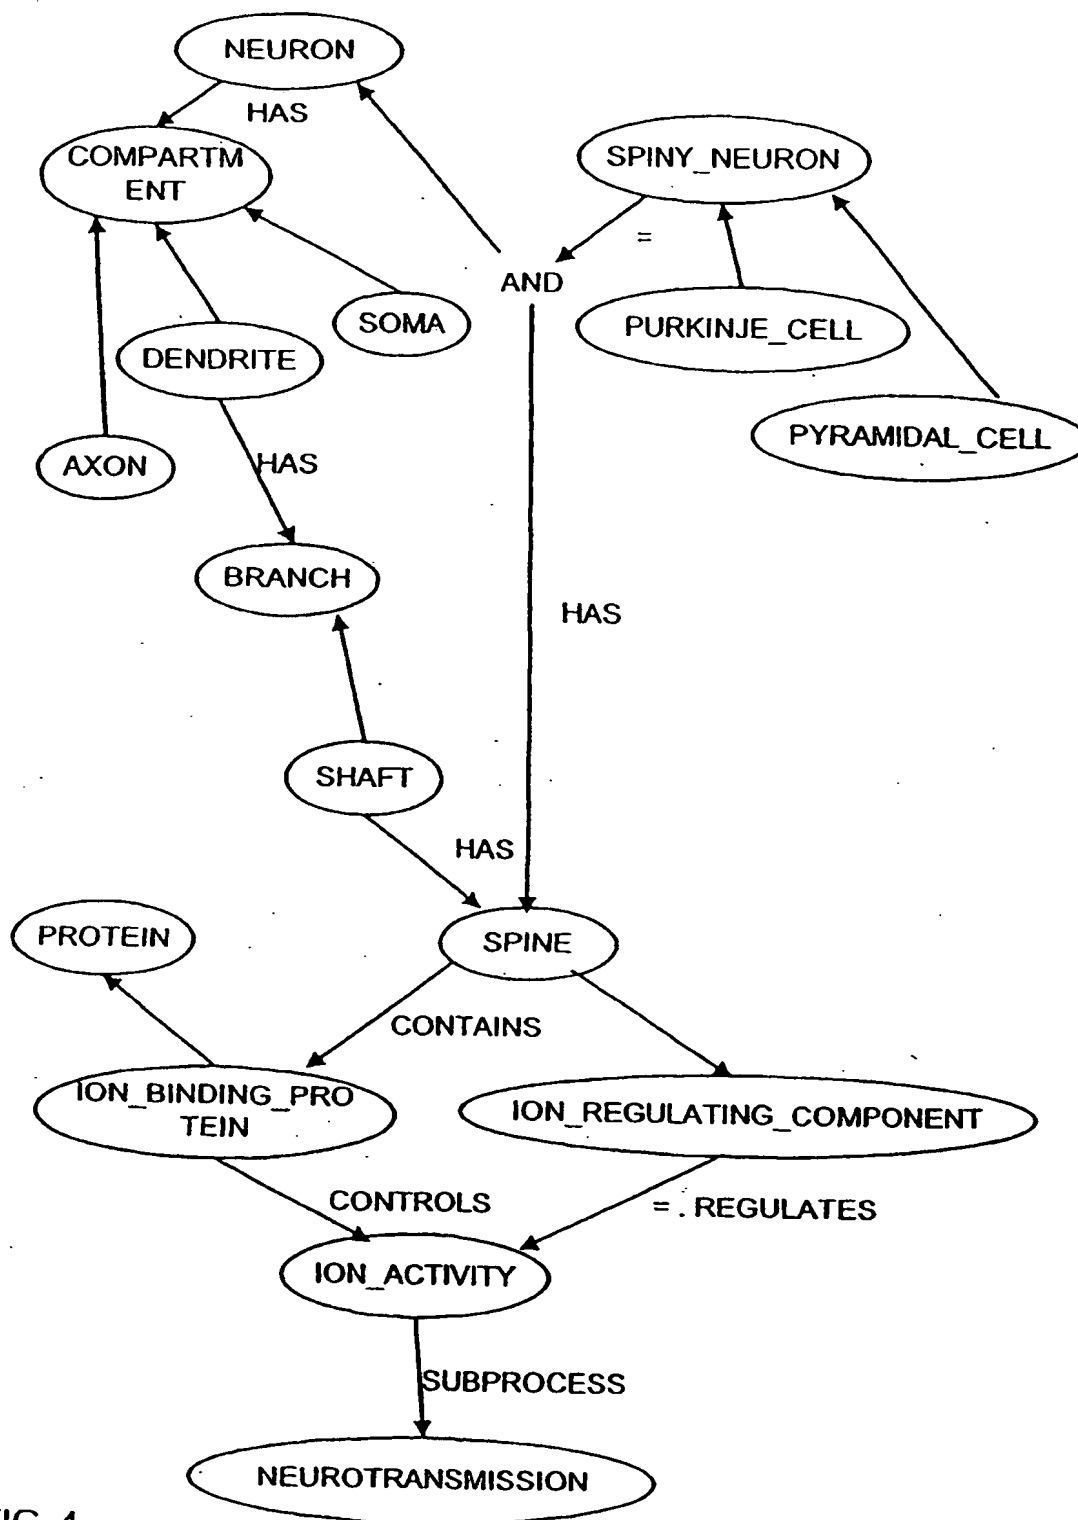


FIG. 4

4/7

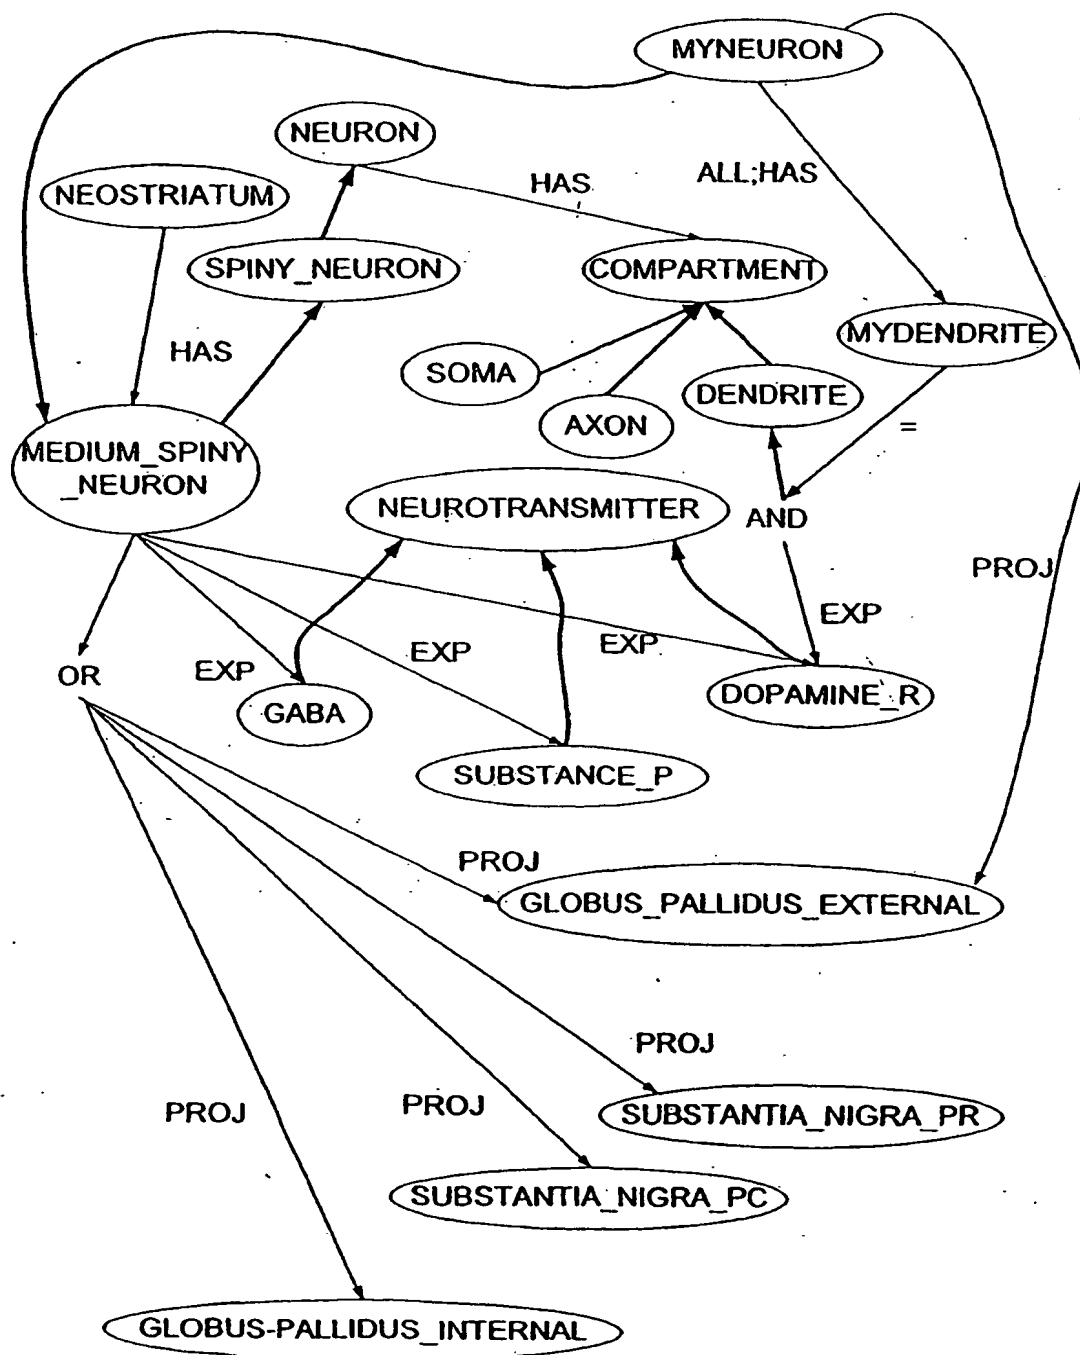


FIG. 5

5/7

PROTLOC Results	
ANATOMY - <i>hypothalamic_pituitary</i>	
PROTEIN - <i>hypothalamic_pituitary</i> (LOCUS: <i>LOCUS</i>) SPECIES - <i>rat</i> IMAGE_ID - <i>02</i> PROTEIN - <i>hypothalamic_pituitary</i> (LOCUS: <i>LOCUS</i>) SPECIES - <i>rat</i> IMAGE_ID - <i>02</i> PROTEIN - <i>hypothalamic_pituitary</i> (LOCUS: <i>LOCUS</i>) SPECIES - <i>rat</i> IMAGE_ID - <i>02</i> PROTEIN - <i>hypothalamic_pituitary</i> (LOCUS: <i>LOCUS</i>) SPECIES - <i>rat</i> IMAGE_ID - <i>02</i>	
PROTEIN - <i>hypothalamic_pituitary</i> (LOCUS: <i>LOCUS</i>) SPECIES - <i>rat</i> IMAGE_ID - <i>02</i> PROTEIN - <i>hypothalamic_pituitary</i> (LOCUS: <i>LOCUS</i>) SPECIES - <i>rat</i> IMAGE_ID - <i>02</i> PROTEIN - <i>hypothalamic_pituitary</i> (LOCUS: <i>LOCUS</i>) SPECIES - <i>rat</i> IMAGE_ID - <i>02</i>	
PROTEIN - <i>hypothalamic_pituitary</i> (LOCUS: <i>LOCUS</i>) SPECIES - <i>rat</i> IMAGE_ID - <i>02</i> PROTEIN - <i>hypothalamic_pituitary</i> (LOCUS: <i>LOCUS</i>) SPECIES - <i>rat</i> IMAGE_ID - <i>02</i> PROTEIN - <i>hypothalamic_pituitary</i> (LOCUS: <i>LOCUS</i>) SPECIES - <i>rat</i> IMAGE_ID - <i>02</i>	

FIG. 8

1. Interactive PROTEIN Localization Query
(UCSD/NCBI data only)

Find the distribution of PROTEIN having ISOFORMS in SPECIES and corresponding IMAGE_ID

SPECIES =

PROTEIN =

ISOFORM =

FIG. 6

6/7

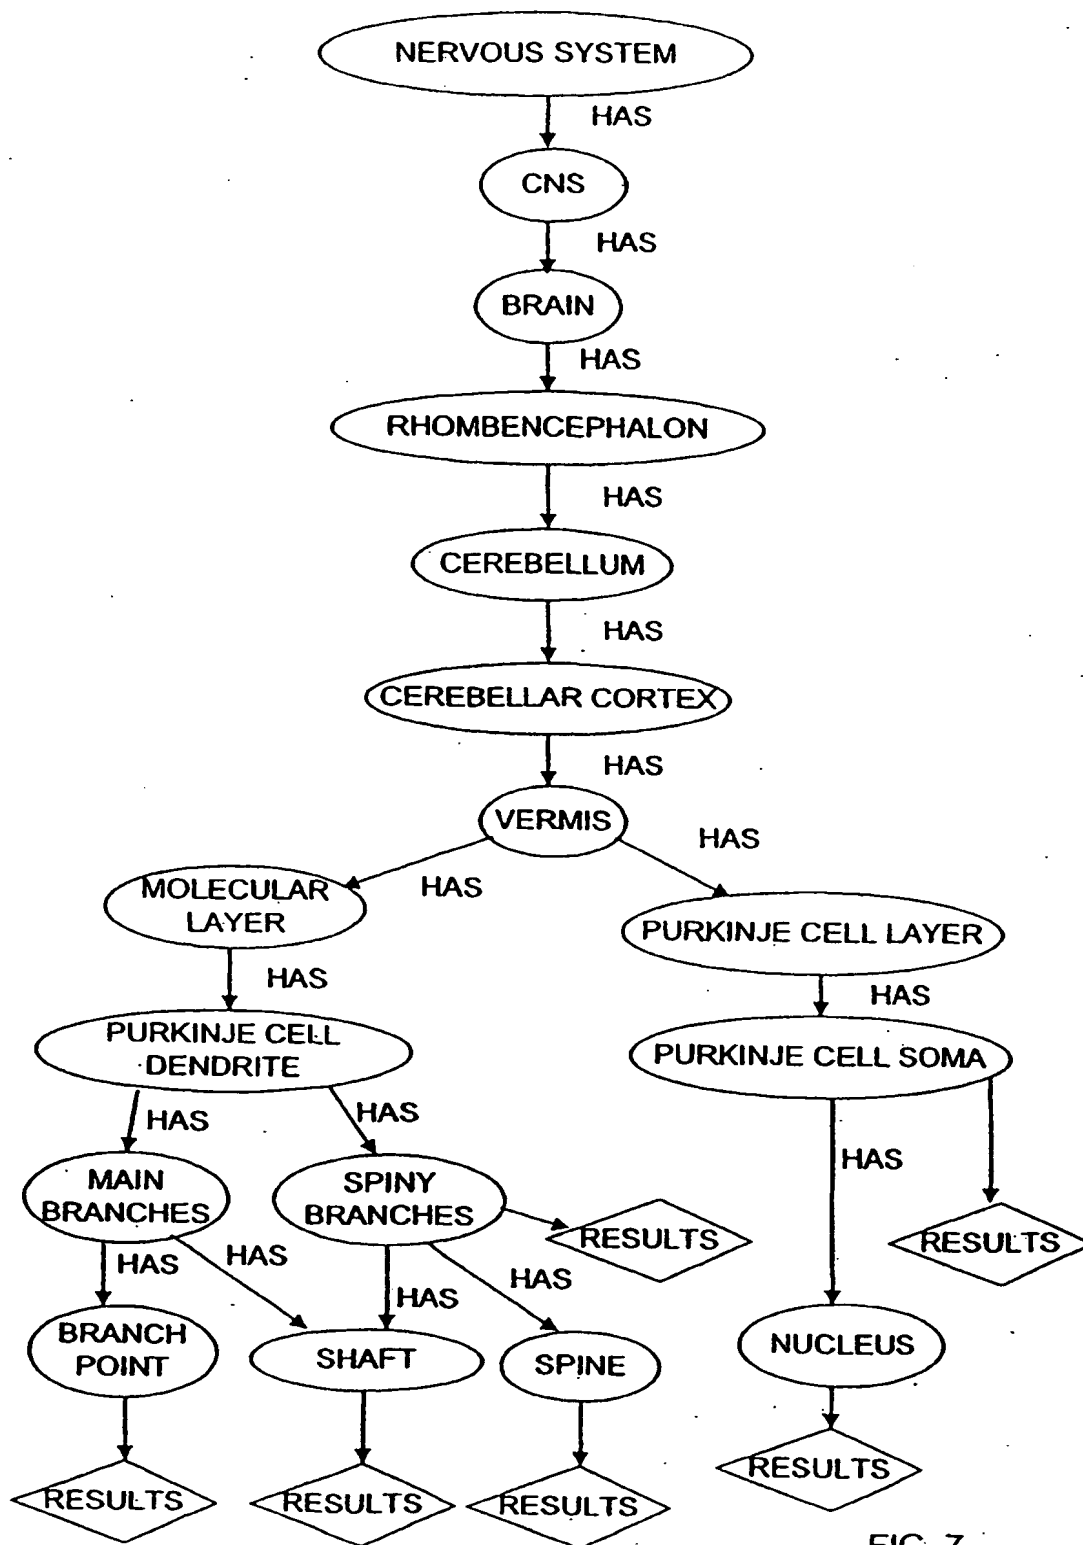


FIG. 7

7/7

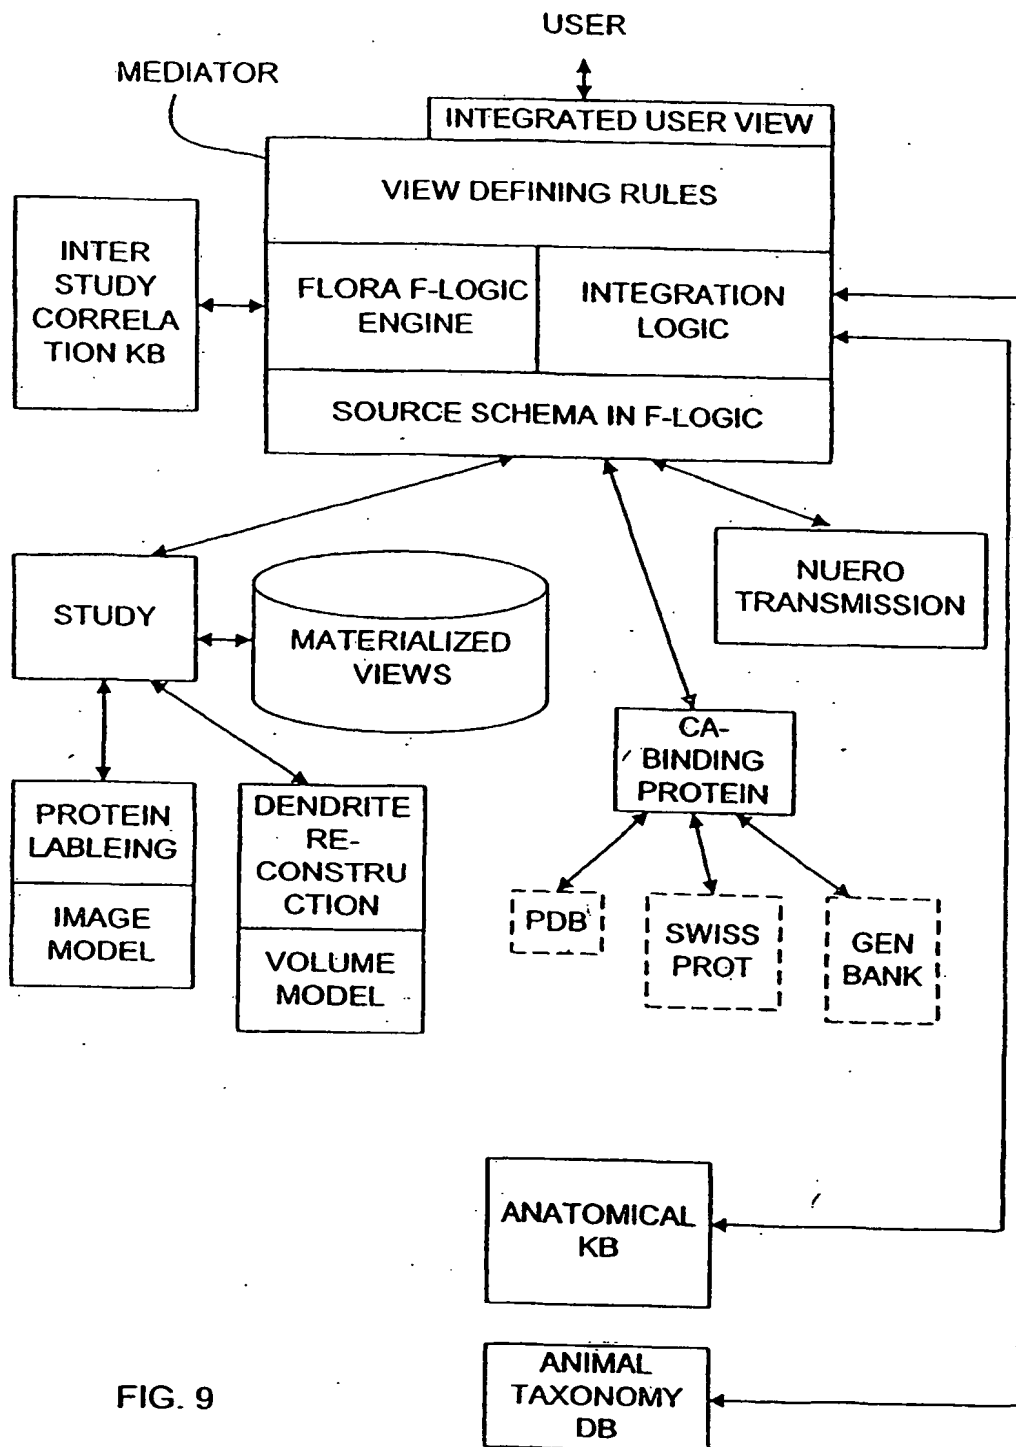


FIG. 9

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US01/27655

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 7/60, 17/10, 101:00

US CL : 703/2

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 703/2

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EAST, ACM, IEEE

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y,P	US 6,134,559 (BRUMME et al.) 17 October 2000, abstract, figures 1, 3, 5, 13-14, 16 and corresponding text	1-48
Y	US 5,995,961 (LEVY et al.) 30 November 1999, abstract, figure 1 and corresponding text.	1-48
Y	US 5,995,959 (FRIEDMAN et al.) 30 November 1999, abstract, figures 1-2 and corresponding text	1-48

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	* Later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"B" earlier document published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Z" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

16 DECEMBER 2001

Date of mailing of the international search report

14 JAN 2002

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3030

Authorized officer

HUGH JONES

Telephone No. (703) 305-9704